PATENT
Docket No. SMI/0005.00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Murray Kucherawy	Examiner: Swearingen, Jeffrey R
Serial No.: 09/945,130	Art Unit: 2145
Filed: August 31, 2001	APPEAL BRIEF
For: E-mail System Providing Filtering Methodology on a Per-Domain Basis	

Mail Stop Appeal
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

BRIEF ON BEHALF OF MURRAY KUCHERAWY.

This is an appeal from the Final Rejection mailed May 25, 2005, in which currently-pending claims 1-53 stand finally rejected. Appellant filed a Notice of Appeal on August 29, 2005 (as indicated by return of a confirmation postcard marked "OIPE AUG 29 2005"). This brief is submitted in triplicate in support of Appellant's appeal.

12/01/2005 MBINAS 00000010 501370 09945130
01 FC:2402 250.00 DA

TABLE OF CONTENTS

1. REAL PARTY IN INTEREST

The real party in interest is assignee Sendmail, Inc., located at 6425 Christie Ave., 4th Floor, Emeryville, CA 94608.

2. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences known to Appellant, the Appellant's legal representative, or assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

3. STATUS OF CLAIMS

Claims 1-53 are pending in the subject application and are the subject of this appeal. An appendix setting forth the claims involved in the appeal is included as the last section of this brief.

4. STATUS OF AMENDMENTS

Two Amendments have been filed in this case. Appellant mailed an Amendment on April 1, 2005, in response to a non-final Office Action dated December 1, 2004. Net of the Amendment, Appellant believes that the pending claims clearly distinguished the claimed invention over the art of record. In response to the Examiner's Final Rejection dated May 25, 2005, Appellant filed a Notice of Appeal. Appellant has filed an Amendment After Appeal to remove non-art issues (as discussed below in section 6). Based on telephone discussion with the Examiner, it is understood that this Amendment (which makes the single text substitution of "desired" for "user-configurable" in the claims) will be entered. Appellant has chosen to forgo any further amendments that might limit is the scope of Appellant's claims, as it is believed that further amendments to the claims are not warranted in view of the art.

5. SUMMARY OF INVENTION

In accordance with the present invention, operation of an e-mail system is

modified to incorporate a flow control filter (service). During processing of incoming e-mail, each child MTA process (that is created to handle a particular new connection) connects to the flow control filter service, so that it can interact with the service during arrival of a message. (See, e.g., Appellant's Specification, p. 17, lines 20-22) This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). (See, e.g., Appellant's Specification, p. 17, lines 22-26) Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. (See, e.g., Appellant's Specification, p. 17, lines 26-29) Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail. (See, e.g., Appellant's Specification, p. 17, line 29 to p. 18, line 2) Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered. (See, e.g., Appellant's Specification, p. 18, lines 2-5)

The overall methodology of operation may be summarized as follows. The following method steps occur in the context of an incoming message that is being processed by the e-mail system (i.e., MTA forking has already occurred) and now the system is ready to evoke the services of the flow control filter of the present invention. (See, e.g., Appellant's Specification, p. 24, line 36 to p. 25, line 1) Invocation of the flow control filter begins with the MTA (i.e., a child MTA of the original (parent) listener) connecting to the flow control filter (e.g., using Sendmail Milter protocol); the filter accepts the connection. (See, e.g., Appellant's Specification, p. 25, lines 3-5; Fig. 6A at 601) The MTA and the filter perform a handshake sequence, including feature and parameter negotiation. At the conclusion of the handshake sequence, a new thread is

created (i.e., in the flow control engine) for processing the new connection/message. (See, e.g., Appellant's Specification, p. 25, lines 5-9; Fig. 6A at 602) Now, the MTA passes to the filter the corresponding connection information (e.g., IP address and host name) of the sending MTA. (See, e.g., Appellant's Specification, p. 25, lines 10-13; Fig. 6A at 603) Based on the connection information, the filter may look up matching class data from the configuration file. (See, e.g., Appellant's Specification, p. 25, lines 13-14; Fig. 6A at 604) In the event that no matching class data is found, the filter will assume unrestricted access for the host and therefore will accept the connection and message. In that case, the flow control engine thread handling the connection may terminate, as there is no further filtering work to be done for this incoming connection and message; the MTA proceeds normally with no further interaction with the filter. (See, e.g., Appellant's Specification, p. 25, lines 14-19; Fig. 6A at 605) Otherwise, the method proceeds to the following filtering steps. The method tests whether class limits have been reached. (See, e.g., Appellant's Specification, p. 25, line 21; Fig. 6A at 606) In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current connection count. (See, e.g., Appellant's Specification, p. 25, lines 21-23; Fig. 6B at 607) Otherwise (i.e., false), the method terminates with the filter rejecting the connection and returning an administrator-defined error code. (See, e.g., Appellant's Specification, p. 25, lines 23-25; Fig. 6B at 608) In the event that the process did not terminate, the MTA reports the sender information to the filter; this occurs in response to the MAIL FROM SMTP phase. (See, e.g., Appellant's Specification, p. 25, lines 25-27; Fig. 6B at 609)

The method notes the sender (i.e., who is the sender) in the class. The administrator-defined class may include, for example, a sender-based parameter indicating that the filter should note the number of unique senders that have arrived in a given timeframe for this particular host (of the class). (See, e.g., Appellant's Specification, p. 25, line 25 to p. 26, line 3; Fig. 6B at 610) In a manner similar to above, the method tests whether class' sender limits have been reached. (See, e.g., Appellant's Specification, p. 26, lines 3-4; Fig. 6B at 611) In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current

unique sender totals. (See, e.g., Appellant's Specification, p. 26, lines 4-6; Fig. 6B at 612) Otherwise, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 6-8; Fig. 6B at 613) In the event that the filtering process has not terminated based on sender information, the method proceeds to test recipient (RCPT TO) information. The configuration file allows the administrator to define a class that limits the number of unique recipients received for that class, over any given time span. As a given message may have multiple recipients, the step repeats for each recipient (information) of the message. (See, e.g., Appellant's Specification, p. 26, lines 9-13; Fig. 6C at 614) As before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 13-15; Fig. 6C at 615) Otherwise, the method updates the totals and proceeds. (See, e.g., Appellant's Specification, p. 26, line 16; Fig. 6C at 616)

The MTA reports the message body, which may be transmitted as one or more blocks. (See, e.g., Appellant's Specification, p. 26, lines 17-18; Fig. 6C at 617) The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. (See, e.g., Appellant's Specification, p. 26, lines 18-20; Fig. 6C at 618) The MTA reports end of message for the current incoming message. (See, e.g., Appellant's Specification, p. 26, lines 20-21; Fig. 6C at 619) The method compares the message size against class limits specified in the configuration file. (See, e.g., Appellant's Specification, p. 26, lines 21-22; Fig. 6D at 620) Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 22-24; Fig. 6D at 621) Otherwise, the incoming message has passed all filters and is accepted. Now, the method may repeat for other incoming messages. (See, e.g., Appellant's Specification, p. 26, lines 24-26; Fig. 6D at 623)

This approach may be easily scaled, for application on a site-wide basis. In that instance, the flow control filter service monitors the children processes for a number of e-mail servers at a given site. In such a configuration, the flow control filter service would

apply policy on a global (site) basis, instead of on a per server basis. (See, e.g., Appellant's Specification, p. 18, lines 6-9)

6. ISSUES

The issues presented on appeal are:

(1) whether claims 1, 16, 17, 20- 22, 25, 27, 28, 31-34, 41, 47, 48, and 50-53 are unpatentable under 35 U.S.C. 102(e);

(2) whether claims 2-4, 5, 9-11, 18, 19, 23, 24, 26, 35, 37, 38, 42, 44 are unpatentable under 35 U.S.C. 103(a);

(3) whether claims 6, 12, 14, 29, 30, 46 are unpatentable under 35 U.S.C. 103(a); and

(4) whether claims 7, 8, 10, 13, 15, 36, 39, 40, 43, 45, 49 are unpatentable under 35 U.S.C. 103(a).

(Duplication across groups (e.g., Claim 10) is necessitated by the Examiner's rejections.)

Regarding the Examiner's non-art rejections in the Examiner's Final Action Paragraph 1 ("Information Disclosure Statement" rejection) and Paragraph 2 (indefiniteness rejection), it is Appellant's understanding that these rejections are overcome by an Information Disclosure Statement filed October 27, 2005, and by Appellant's Amendment After Appeal filed October 31, 2005.

7. GROUPING OF CLAIMS

For purposes of this appeal, Appellant believes that the following groups of claims are separately patentable under Sections 102 and 103. Thus, the claims do not stand or fall together with respect to the rejections under Sections 102 and 103 but are instead grouped as follows:

Group	Claims
1	1, 16, 17, 20-22, 25, 27, 28, 31-34, 41, 47, 48, 50-53
2	2-4, 5, 9-11, 18, 19, 23, 24, 26, 35, 37, 38, 42, 44
3	6, 12, 14, 29, 30, 46
4	7, 8, 10, 13, 15, 36, 39, 40, 43, 45, 49

(The reasoning supporting separate patentability of the above groups is set forth in detail below, in the Argument section.)

8. ARGUMENT

A. Rejection under Section 102

1. General

A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in the single prior art reference. (See, e.g., MPEP Section 2131.) As will be shown below, the reference fails to teach each and every element set forth in claim 1, as well as other claims of Group 1, and therefore fails to establish anticipation of the claimed invention under Section 102.

2. Group 1 claims: Srivastava et al.

Claims 1, 16-17, 21-22, 25, 27-28, 31-34, 41, 47-48 and 50-53 (and apparently claim 20) stand rejected under 35 U.S.C. 102(e) as being anticipated by Srivastava et al. (U.S. Patent No. 6,374,292), hereinafter referred to "Srivastava"). The Examiner's rejection of claim 1 is representative:

> Regarding claim 1, Srivastava discloses a method for processing an incoming e-mail message that is being received from another domain, the method comprising: receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system; in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis; comparing the request from the particular domain against configurable policy rules; and denying the request if any of said policy rules would be violated. [Srivastava discloses using a process to define a particular domain in an email server. An individual

> domain can be configured to allow all mail to be received if the state of the domain is active (establish a new connection) or if the state of the domain is inactive the domain is suspended from routing mail (denying the request). See Srivastava, column 7, lines 36-59. Srivastava further discusses using a multithreaded process, with each thread handling a connection. Examiner considers this to be equivalent to creating a second process for handling a new connection. Srivastava further states that using a single multithreaded process is beneficial by maximizing performance and stability and by minimizing system resource usage. See Srivastava, column 5, lines 9-15.] By this rationale claim 1 is rejected.

As shown below, Appellant's claimed invention is distinguishable on a variety of grounds.

Srivastava refers to a package of software running on a mail server which governs which specific e-mail related services are offered to users in a particular domain or set of domains. In fact, all of the services described in Srivastava are at the upper OSI layers, i.e., "Presentation" and "Application". Chief among these is virtual domain hosting, whereby one server is given the ability to send and receive e-mail for a variety of otherwise unrelated domains, and to thereby provide e-mail services for users within those domains. The intent of Srivastava is to give the ISP (Internet Service Provider) the means to delegate these services to the administrators of those respective domains without also maintaining individual e-mail servers for each domain or granting machine-wide administrative access to lots of disparate organizations, which of course are prospects with very serious scaling implications.

Appellant's flow control invention, since it operates below the "Presentation" layer in the OSI model, has little knowledge of which domains are stored on the server(s) it protects. It is instead operating in the "Session" layer, with a small amount of information made available to it from lower layers. Moreover it has no knowledge of any protocol or service related to e-mail other than SMTP. It may even not be running on a server providing SMTP service at all. Instead, it has knowledge of the origin of the TCP connection being made to an SMTP service, and makes use of this knowledge to classify the connection and thereby moderate the flow of e-mail into or out of a system. The foregoing is provided to afford the reader context for understanding basic distinctions between the two approaches. While Srivastava's invention is more geared toward

assistance and delegation of administration of e-mail services, Appellant's flow control invention serves to prevent domination of network and system resources by a particular e-mail source. This core distinction will now be reviewed in further detail, with emphasis given to Srivastava's actual teachings, Appellant's specification, and Appellant's specific claim limitations.

Appellant's flow control invention is designed to moderate the flow of e-mail inbound. An example of this might be to limit the maximum number of e-mails, connections over time, or simultaneous connections coming from a **given domain** (say, e.g., "prodigy.net") so that the impact of a sudden surge in spam or a deliberate denial-of-service attack can be detected and quashed without impacting the flow of e-mail from other domains. An outbound policy can also be applied, so that for example a user's desktop infected with a virus that then tries to send e-mail to hundreds or thousands of users is blocked when it is detected. Both of these concepts are known commonly as "traffic shaping", although that term often describes a function implemented by routers at the OSI "Transport" layer and below. This is in contrast to Srivastava, which is essentially an e-mail server provisioning system, allowing a machine or set of machines to be "sliced up" in such a way that many "virtual" domains can be served by a small number of hosts (or one), and the services provided to those domains can be managed by the domain's respective owners.

Turning now to the Examiner's specific basis for rejection, the Examiner analogizes Appellant's flow control filter to Srivastava's virtual domains (Srivastava, Col. 7, lines 36-59).

> Referring now to FIG. 4, showing a flowchart that details a process 500 for defining a virtual domain in accordance with an embodiment of the invention. The process 500 begins at 502 by defining a virtual domain node in the DIT. Once the virtual domain node has been defined, corresponding routing table entries are defined at 504 and at 506, various virtual domain are stored at the virtual domain node. It should be noted that the various virtual domain include a list of services permitted the domain. Such services include IMAP, MAPS, POP3, POP3S, SMTP which in some cases requires presentation of credentials. Other of the services include identification of a domain administrator who is authorized to manage the particular virtual domain

> which includes setting particular user-level for particular users in the domain. These services also include designation of a virtual domain postmaster who identifies email message delivery problems, and a state of the domain.

As clearly shown above, Srivastava at this point is describing the creation of a virtual domain. The purpose of a "virtual domain" is discussed by Srivastava: "It is therefore desirable that an email service provider be able to offer email services to multiple organizations each of which has their own virtual domain and to support the ability to define such domains in the directory and host them on a shared mail server." This is not the same as Appellant's limitation of "said second process being connected to a flow control filter providing filtering on a per-domain basis," which is able to block inbound e-mail traffic -- or allow e-mail traffic -- from different domains, depending on whether a particular given domain is complying with the "configurable policy rules."

For example, if a given domain in Srivastava's system is specified to be a virtual domain for which e-mail services are allowed, then Srivastava's system would permit all e-mail traffic from that virtual domain -- regardless of whether some of that traffic is coming from a user machine engaged in spam or a deliberate denial-of-service attack. There certainly is no mention or passing suggestion in Srivastava that his system also includes some sort of adaptive filter that would then somehow further monitor the virtual domains as they make connections to determine whether the connections for e-mail traffic originating from those domains violate policies in a manner that would cause his system to begin rejecting e-mail traffic until such time as the noncompliant domain returns to compliance. Thus, a detailed review of Srivastava's disclosure that the Examiner relies on for the rejection reveals that it is entirely silent regarding any feature that could function in a manner that is analogous to Appellant's claimed flow control mechanism. In order to achieve Appellant's result (e.g., blocking spam and denial of service attacks) in Srivastava's system, one would have to add Appellant's filtering mechanism to Srivastava's system.

Further, the Examiner points to Srivastava's Col. 5, lines 9-15, which states:

> In the described embodiment, access to the message store 304 is multithreaded thereby allowing a single process to manage a large number of connections since each connection is handled by a thread. In this way, multithreaded access maximizes both performance and scalability by minimizing the system resources required for the management of each connection.

Here, the Examiner contends that Srivastava's multithreaded access to a single message store is the same as Appellant's approach of spawning a second process for handling a new incoming connection.

"Threads" and "processes" are not the same. A "process" is an executing program or task. A "thread" is a part of a process that can execute independently of other parts; it exists within a process and uses the process' resources. Unlike processes, multiple threads run within the same address space and share their process' data. The concepts of threads and processes are well known and well documented in the technical literature. (A copy of Kalev, Danny, "Processes and Threads," ITWorld.com, February 9, 2001, was attached for the Examiner's convenience to Appellant's first-filed Amendment.) The article discussed threads and processes in the context of the Linux operating system, but the discussed concepts apply equally well to other operating systems (e.g., UNIX, Windows, Macintosh OS X).

Without discussing the other deficiencies of Srivastava at this point (e.g., a single "message store" in Srivastava versus multiple incoming connections from different domains in Appellant's system), it is clear that the section that the Examiner cites in Srivastava discusses the use of multiple threads, not the spawning of additional processes. If anything, Srivastava at this point teaches away from Appellant's claimed approach of multiple processes (not Srivastava's approach of a single process with multiple threads).

Appellant's flow control invention provides a facility for moderating the flow of SMTP traffic (connections, aggregate volume, and unique senders) into a server or set of servers. This feature is brought out in Appellant's claims. For example, Appellant's claim 1 recites:

> receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;
>
> (Emphasis added.)

This is an incoming connection from another domain (particularly, an MTA at another domain), for the purpose of doing an MTA-to-MTA email exchange. This is not an operation for servicing user requests.

Further, claim 1 requires:

> in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;
>
> comparing the request from the particular domain against configurable policy rules; and
>
> denying the request if any of said policy rules would be violated.
>
> (Emphasis added.)

Here, the claim requires that the above-mentioned incoming connection is passed through a **domain-specific** filter. This approach allows Appellant's flow control invention to detect and prevent massive spam from being received on incoming connections of a particular domain. Srivastava's approach of granting user services cannot be morphed into system that prevents incoming massive spam from an MTA of a particular domain.

Srivastava is essentially a method for providing virtual hosting services for e-mail and web pages, with the ability to create virtual users within that context and optionally delegate authority to those users to manage parts of the virtual space so provisioned. On startup, Appellant's flow control invention reads a configuration file and then reacts to SMTP traffic it observes. It has no "user-serviceable parts"; only the e-mail administrator has access to read or change its configuration. Although the two approaches converge insofar as they are both related *generally* to providing e-mail service at ISPs and can do

some amount of per-user validity checking, the convergence ends there. They otherwise operate on different types of data and operate in different manners (and in fact in different layers of the OSI protocol stack). By virtue of these core architectural differences, Appellant's claims include very specific claim limitations (explicitly highlighted above) that are not taught or suggested by Srivastava. It is respectfully submitted that these features, as set forth in Appellant's claims, clearly distinguish over Srivastava.

The other independent claims rejected under Section 102 (i.e., claims 21 and 41) include the above-mentioned distinguishing per-domain filtering or policy enforcement claim limitations, and are therefore believed to be allowable for the reasons stated above. (The dependent claims rejected under Section 102 are believed to be allowable by virtue of depending from the foregoing independent claims.) Accordingly, it is respectfully submitted that the claims distinguish over Srivastava, and that the Examiner's rejection under Section 102 should not be sustained.

B. Rejections under Section 103

1. General

Under Section 103(a), a patent may not be obtained if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. To establish a prima facie case of obviousness under this section, the Examiner must establish: (1) that there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings, (2) that there is a reasonable expectation of success, and (3) that the prior art reference (or references when combined) must teach or suggest all the claim limitations. (See e.g., MPEP 2142). The references cited by the Examiner fail to meet these conditions.

2. Group 2 claims: Srivastava combined (in various permutations) with Apache, Mosberger, Ahmed, Shaw, Rakoshitz, and Bates

Claim 35 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and RFC 821. Claim 26 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Apache HTTP Server Configuration Files ("Apache"). Claims 2 and 42 (and apparently claim 3) stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Mosberger et al. (U.S. Patent 6,438,597, "Mosberger"). Claims 18, 19, 23, and 24 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Ahmed et al. (U.S. Patent No. 6,704,772). Claim 9 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw et al. (U.S. Patent No. 6,282,565, "Shaw"). Claims 11 and 44 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of Sash (U.S. Pub. No. 2003/0167250). Claims 4 and 5 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Rakoshitz et al. (U.S. Patent No. 6,816,903, "Rakoshitz"). Claim 10, 37, and 38 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of RFC 821.

For the rejected claims of this group, the Examiner has essentially repeated his Srivastava rejection (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section), but has combined bits and pieces of other art references in an effort to shore up his base Srivastava rejection. The various combinations fail to establish a prima facie rejection under Section 103, as the combined references fail to teach or suggest all the claim limitations of the claims of this group. Importantly, the deficiencies of the base Srivastava rejection are left wholly unaddressed.

For example, Mosberger describes firewall-like features of controlling connections (e.g., based on domain). However, Mosberger does not provide sufficient teaching to morph Srivastava's virtual domain hosting into an e-mail flow control filter that controls connections based on domain-specific behavior of e-mail traffic. As another example, Rakoshitz describes "select counters for monitoring incoming and outgoing traffic from a link" (Rakoshitz, at Col. 21, lines 2-3). Nowhere does Rakoshitz describe maintaining "a counter indicating how many connections have been granted to the particular domain" (emphasis added), as required by Appellant's claim 4, for example. To

the extent that Rakoshitz teaches a counter, the Rakoshitz counter is one that tracks individual links. In particular, no description is given which teaches or suggests that the individual links traceable to or referencing a particular domain be tracked with a counter. Accordingly, at best, Rakoshitz teaches away from Appellant's domain counter claim limitation.

As yet another example, the Examiner adds the Sash permutation/combination for the additional teaching of "a maximum number of different recipients permitted..." Sash describes an information template and describes limiting a maximum number of recipients that an information template can be sent to (i.e., limit the number of times it can be forwarded to other recipients). Appellant's claim limitation states, "said configurable policy rules specify a maximum number of different recipients **permitted by a given domain** over a user-configurable period of time." (See, e.g., Appellant's claim 11.) This would apply, for instance, in this scenario of e-mail traffic coming from a particular domain (e.g., *advertiser.net*) having an inordinate number of different recipients (say, e.g., > 10 million). Placing a restriction on the number of times that a data object can be forwarded (e.g., Sash's restriction on the number of recipients that Sash's information template can be forwarded to) bears little relevance to Appellant's claim limitation.

Without even getting to the issue of whether there is some suggestion or motivation in these references to make the combination suggested by the Examiner, the numerious art rejection permutations that the Examiner has put together for the claims of this group all fails to meet even the most basic threshold of teaching or suggesting all the claim limitations. Importantly, the incremental art references added by the Examiner to create the four Section 103 rejections for the claims of this group each fail to cure the deficiencies of Srivastava regarding the core elements set forth in Appellant's base claims (which the present claims depend from). It is respectfully submitted that the various rejections of claims of this group have each failed to establish a prima facie case of obviousness under Section 103, and therefore it is respectfully requested that the Examiner's rejections of these claims not be sustained.

3. Group 3 claims: Srivastava combined with Spam!

Claims 6, 12, 14, 29, 30 and 46 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! (Cranor and LaMacchia, Communications of the ACM, August 1998). Here, the Examiner relies on Srivastava as above, and adds Spam! for the proposition that it is obvious to add a spam filter. The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava in its failure to teach or suggest Appellant's per-domain flow control filter (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section). Besides not curing the deficiencies of Srivastava, the Spam! reference, as used by the Examiner for rejecting claims of this group, is particularly deficient and thus warrants additional discussion and consideration.

At the outset, it is important to recognize that Appellant does not claim to be the first to invent a spam filter, and Appellant's invention itself is not a spam filter but instead a flow control filter which operates at e-mail server level (message transport agent or "MTA") to monitor the behavior of different domains that are connecting to send incoming e-mail. For example, a "bad" behavior would include a denial-of-service attack, which of course itself is not "spam" (unsolicited e-mail) as that term is generally understood. And, of course, a "spam" filter (such as discussed by the Spam! reference) would be useless against such non-spam malicious behavior. This core distinction will now be examined in further detail, with particular emphasis on Appellant's claim limitations.

With respect to claim 6, for example, Appellant's claim discusses "comparing the sender information **from the particular domain** against said configurable policy rules." This is not simply rejecting an e-mail piece based on it coming from a specific sender that has been blacklisted. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following teaching from Appellant's specification (at page 17, line 20 to page 18, line 5):

> As described above, with each new connection a child MTA process is created. In accordance with the present invention, each child process connects to the flow control filter service, so that it can interact with the

> service during arrival of a message. This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. **By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail encountered for a particular domain over a given period of time.** Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered.
>
> (Emphasis added.)

As described above, the purpose of examining header information is not to accept or reject a given single piece of e-mail based on spam criteria (e.g., blacklisted or not), but is use in conjunction with Appellant's flow control filter to further characterize the given domain that is being monitored. For example, as specified in the passage above, one of the criteria may be "number of different recipients." This requires the system to look at e-mail header information, but note in particular that the header information is not being examined for purposes of identifying an individual piece of e-mail as spam, but instead is being used to further characterize the current behavior of the given **domain** that is being monitored (in order to determine whether server-level intervention is warranted). Should anyone have any doubt as to the intention of Appellant's claim language, one need only refer back to parent claim 1, which clearly describes Appellant's invention as a "flow control filter providing filtering on a **per-domain** basis."

Claim 12 similarly establishes that Appellant's claim limitations are not simply claiming or reading on a spam feature. Appellant's claim 12 discusses comparing the e-

mail message body data against policy rules for a particular domain. This is not simply rejecting an e-mail piece based on it having certain content (e.g., explicit content) that is detected and rejected by a spam filter. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following from Appellant's specification (at page 10, line 27 to page 11, line 2):

> The MTA reports the message body, which may be transmitted as one or more blocks. **The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. The MTA reports end of message for the current incoming message. The method compares the message size against class limits specified in the configuration file.** Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code).
>
> (Emphasis added.)

As described above, the message header may be examined, not for the purpose out of accepting or rejecting a given single piece of e-mail based on spam criteria (e.g., offensive content), but is used in conjunction with Appellant's flow control filter to further characterize the given source -- a particular domain (as required by claim 12, and base claim 1) -- that is being monitored. For example, as specified in the passage above, one of the criteria may be "aggregate total bytes received from a given source over a period of time." (See, e.g., Appellant's claim 46: "determining a maximum aggregate volume of e-mail permitted for the particular domain over a user-configurable period of time.") This requires the system to look at the message body, but note in particular that the message body is not being examined for purposes of identifying the e-mail as having spam content, but instead is being used to further characterize the current behavior of the given domain that is being monitored.

Regarding claim 14, for example, the Examiner contends that this is taught by Spam! at page 79, which indicates that ISPs may limit the number of outbound messages each subscriber can send. However, that is not Appellant's claim limitation. Instead,

claim 14 recites (net of amendments) that "said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a user-configurable period of time." As readily apparent from the claim language, the volume of e-mail being regulated is that coming from a given domain, not that coming from an individual user or subscriber. Thus, for example, applying Appellant's invention, the *uspto.gov* e-mail server could be configured to detect an abnormal amount of e-mail coming from *aol.com* (and take appropriate action, accordingly), for example as a result of a denial-of-service attack originating from that domain. Such a result cannot be achieved by simply using a spam filter at the ISP (*aol.com*) for attempting to detect an abnormally high level of e-mail from any given subscriber. And, in the case of a distributed denial-of-service attack, the attack may be distributed over numerous subscriber accounts (e.g., as a result of Trojan/zombie infection), and it may very well be the case that no one subscriber has an abnormal level of outbound messages.

Regarding claims 29 and 30, for example, the Examiner contends that Spam! at p. 79 teaches that limits can be placed on a domain. However, Spam! itself describes placing limits on individual subscriber accounts. Spam! only describes blocking e-mail from a bogus domain. It contains no description of how a spam filter could continually monitor the e-mail traffic behavior of a given domain, and apply configurable policy rules. Again, the characteristics of e-mail traffic coming from a given domain (e.g., *aol.com*) are not the same as the characteristics of e-mail traffic coming from an individual subscriber (e.g., *john.smith@aol.com*).

For the reasons stated above, the references cited by the Examiner fail to teach or suggest all the claim limitations. Accordingly, it is respectfully submitted that the claims of this group distinguish over the references and are patentable under Section 103. Thus, it is respectfully requested that the Examiner's rejection not be sustained.

4. Group 4 claim: Srivastava/Spam! further combined (individually) with Shaw Bates, and RFC 821

Claim 15 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 14 above, and further in view of Shaw. Claims

8, 43, 45, and 49 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of Bates et al. (U.S. Patent No. 6,779,021, "Bates"). Claim 10 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 9 above, and further in view of RFC 821. Claims 7, 13, 36, 39, and 40 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of RFC 821.

For the rejected claims of this group, the Examiner has essentially repeated his Srivastava/Spam! rejection (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section), but has combined other stray pieces of art in an effort to shore up his rejection. The various combinations fail to establish any competent prima facie rejection under Section 103, as the combinations each fail to teach or suggest all the claim limitations of the claims of this group. Importantly, the deficiencies of the base Srivastava/Spam! Section 103 rejection (discussed above) are left wholly unaddressed. Notable deficiencies in the Examiner's reading of the additional references is evident, as will now be described.

For example, the Examiner relies on Shaw for disclosing the claim limitation of "limiting the size of incoming e-mail messages based on a maximum number of bytes." (See, e.g,. Examiner's Action, paragraph 81.) However, this is not what Appellant's claim states. Instead, claim 15 recites: "said **maximum aggregate volume** is based on total byte count of e-mail received from a given domain over a user-configurable period of time." (Emphasis added.) Claim 14 (claim 15's parent) recites: "wherein said configurable policy rules specify a maximum aggregate volume of e-mail **permitted by a given domain** over a user-configurable period of time." (Emphasis added.) Limiting the maximum average volume from a given domain is not the same as limiting the size of a given incoming e-mail message. Accordingly, the cited art does not serve as an appropriate rejection.

As another example, the Bates passage cited by the Examiner comes from the Bates' Background Section where Bates describes basic spam filtering techniques, such as blocking on sender. Appellant's claim limitation, however, requires: "a maximum

number of different senders permitted by **a given domain** over a user-configurable period of time." (Emphasis added.) A review of Bates indicates no such feature described or suggested. Further, to the extent that Bates repeats spam filtering information about blocking without regard to the behavior of a particular domain, Bates teaches away from Appellant's claimed approach.

The hodgepodge of art combinations that the Examiner has strung together for the claims of this group each fails to meet even the most basic threshold of teaching or suggesting all the claim limitations. And the incremental art references added by the Examiner not only fail to cure the deficiencies of Srivastava/Spam!, but they have also been misinterpreted or mischaracterized by the Examiner to such an extent that their purported incremental teaching is in fact not present in the references themselves. It is respectfully submitted that the rejections of claims of this group each fail to establish a prima facie case of obviousness under Section 103, and therefore it is respectfully requested that the Examiner's rejection of these claims not be sustained.

9. CONCLUSION

The present invention greatly improves the ease and efficiency of running an e-mail server by providing a domain-based e-mail filter. It is respectfully submitted that the present invention, as set forth in the pending claims, sets forth a patentable advance over the art.

In view of the above, it is respectfully submitted that the Examiner's rejections under 35 U.S.C. Section 102 and 103 should not be sustained. If needed, Appellant's undersigned attorney can be reached at 408 884 1507. For the fee due for this Appeal Brief, please refer to the attached Fee Transmittal Sheet. This Brief is submitted in triplicate.

Respectfully submitted,

Date: November 29, 2005

Digitally signed by John A. Smart
Date: 2005.11.29 18:56:56 -08'00'

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX

10.APPENDIX OF CLAIMS ON APPEAL

1. (Original) In an electronic mail (e-mail) system, a method for processing an incoming e-mail message that is being received from another domain, the method comprising:

receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;

in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;

comparing the request from the particular domain against configurable policy rules; and

denying the request if any of said policy rules would be violated.

2. (Previously presented) The method of claim 1, wherein said configurable policy rules specify a maximum number of connections permitted by a given domain over a user-configurable period of time.

3. (Previously presented) The method of claim 2, wherein said user-configurable period of time is configurable.

4. (Original) The method of claim 1, further comprising:

if none of said policy rules would be violated, permitting the requested connection and incrementing a counter indicating how many connections have been granted to the particular domain.

5. (Previously presented) The method of claim 4, further comprising:

after passage of the user-configurable period of time, resetting the counter.

6. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving sender information about the particular e-mail message from the particular domain;

comparing the sender information from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

7. (Original) The method of claim 6, wherein said sender information is transmitted during a "MAIL FROM" phase of SMTP (Simple Mail Transport Protocol) processing.

8. (Previously presented) The method of claim 6, wherein said configurable policy rules specify a maximum number of different senders permitted by a given domain over a user-configurable period of time.

9. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving recipient information about the particular e-mail message from the particular domain;

comparing the recipient information from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

10. (Previously presented) The method of claim 9, wherein said recipient information is transmitted during a "RCPT TO" phase of SMTP (Simple Mail Transport Protocol) processing.

11. (Previously presented) The method of claim 9, wherein said configurable

policy rules specify a maximum number of different recipients permitted by a given domain over a user-configurable period of time.

12. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving e-mail message body data about the particular e-mail message from the particular domain;

comparing the e-mail message body data from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

13. (Previously presented) The method of claim 12, wherein said e-mail message body data is transmitted during a "DATA" phase of SMTP (Simple Mail Transport Protocol) processing.

14. (Previously presented) The method of claim 12, wherein said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a user-configurable period of time.

15. (Previously presented) The method of claim 14, wherein said maximum aggregate volume is based on total byte count of e-mail received from a given domain over a user-configurable period of time.

16. (Original) The method of claim 1, wherein said first process comprises a mail transport agent (MTA) process.

17. (Original) The method claim 16, wherein said second process comprises a child mail transport agent (MTA) process.

18. (Original) The method of claim 1, wherein said second process is created from said first process via a forking operation.

19. (Original) The method of claim 18, wherein said second process is initially created as a copy of said first process.

20. (Original) The method of claim 1, further comprising:

creating a multitude of new processes for handling multiple requests to establish new connections, each new process being connected to said flow control filter providing filtering on a per-domain basis.

21. (Original) An electronic mail (e-mail) system providing filtering of incoming e-mail messages on a per-domain basis, the system comprising:

a parent process for receiving requests from different domains to establish new connections for transmitting e-mail messages;

a plurality of child processes for handling the requests to establish new connections and for handling subsequent requests for transmitting e-mail messages;

a set of rules specifying conditions for accepting requests for new connections and for accepting requests for transmitting e-mail messages; and

a flow control filter, in communication with said child processes and said set of rules, providing filtering based on each domain's conformance to said rules.

22. (Original) The system of claim 21, wherein said parent process and said child processes comprise mail transport agent (MTA) processes.

23. (Original) The system claim 21, wherein each said child process is created from the parent process via a forking operation.

24. (Original) The system of claim 21, wherein each said child process is initially created as a copy of said parent process.

25. (Original) The system of claim 21, wherein said set of rules comprises a configurable set of rules.

26. (Original) The system of claim 21, wherein said set of rules comprises a set of rules stored in a text-based configuration file.

27. (Original) The system of claim 21, wherein said set of rules comprises user-created class definitions specifying different classes of domains.

28. (Original) The system of claim 27, wherein each said class definition includes a domain name corresponding to a particular domain that is to be monitored for filtering.

29. (Previously presented) The system of claim 27, wherein each said class definition includes limits that a particular domain must adhere to over a given user-configurable period of time.

30. (Original) The system of claim 29, wherein said limits include selected ones of:

maximum number of different senders,
maximum number of different recipients,
maximum number of connections,
maximum number of envelopes, and
maximum aggregate volume of mail.

31. (Original) The system of claim 21, wherein a given domain is not filtered if a corresponding rule has not been created for that given domain.

32. (Original) The system of claim 21, wherein said flow control filter denies a given domain's request for a new connection if any of said rules would be violated by

granting the request.

33. (Original) The system of claim 21, wherein said requests for transmitting e-mail messages comprise SMTP (Simple Mail Transport Protocol) commands submitted to the e-mail system from different domains.

34. (Original) The system of claim 33, wherein said flow control filter processes said SMTP commands received from different domains to ascertain whether any of said rules would be violated.

35. (Original) The system of claim 34, wherein said SMTP commands include a "MAIL FROM" command specifying sender information for a given e-mail message.

36. (Original) The system of claim 35, wherein said flow control filter examines said sender information to ascertain whether any of said rules would be violated.

37. (Previously presented) The system of claim 34, wherein said SMTP commands include a "RCPT TO" command specifying recipient information for a given e-mail message.

38. (Original) The system of claim 37, wherein said flow control filter examines said recipient information to ascertain whether any of said rules would be violated.

39. (Original) The system of claim 34, wherein said SMTP commands include a "DATA" command specifying e-mail message body data for a given e-mail message.

40. (Original) The system of claim 39, wherein said flow control filter examines said e-mail message body data to ascertain whether any of said rules would be violated.

41. (Original) In an electronic mail (e-mail) system, a method for processing

incoming e-mail messages that are being received from different domains, the method comprising:

receiving requests from different domains to establish new connections for transmitting e-mail messages to the e-mail system;

for each request received in connection with transmitting a given e-mail message, performing substeps of:

identifying a particular domain that has submitted the request,

based on the determined identity of the domain, determining whether the request to establish a new connection can be granted without violating policy rules, and

based on the determined identity of the domain, determining whether subsequent requests to transmit different portions of a given e-mail message can be granted without violating said policy rules.

42. (Previously presented) The method of claim 41, wherein said step of determining whether the request to establish a new connection can be granted includes:

determining a maximum number of connections permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of connections if the request were granted.

43. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum number of different senders permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different senders if the request were granted.

44. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail

message can be granted includes:

determining a maximum number of different recipients permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different recipients if the request were granted.

45. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum number of different e-mail envelopes permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different e-mail envelopes if the request were granted.

46. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum aggregate volume of e-mail permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum aggregate volume of e-mail if the request were granted.

47. (Original) The method of claim 41, further comprising:

if the request to establish a new connection cannot be granted without violating said policy rules, denying the request.

48. (Original) The method of claim 47, further comprising:

returning an error code indicating why the request is denied.

49. (Original) The method of claim 41, further comprising:

if the request to transmit different portions of a given e-mail message cannot be granted without violating said policy rules, denying the request.

50. (Original) The method of claim 41, wherein portions of a given e-mail message include sender information, recipient information, and message body data.

51. (Original) The method of claim 41, wherein said policy rules are configurable.

52. (Original) The method of claim 41, wherein said policy rules comprise user-edited rules created for different domains.

53. (Previously presented) The method of claim 52, wherein each user-edited rule comprises a host class definition specifying a particular domain and corresponding limits to be applied against that domain over a user-configurable period of time.

PATENT
Docket No. SMI/0005.00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Murray Kucherawy	Examiner: Swearingen, Jeffrey R
Serial No.: 09/945,130	Art Unit: 2145
Filed: August 31, 2001	APPEAL BRIEF
For: E-mail System Providing Filtering Methodology on a Per-Domain Basis	

Mail Stop Appeal
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

BRIEF ON BEHALF OF MURRAY KUCHERAWY.

This is an appeal from the Final Rejection mailed May 25, 2005, in which currently-pending claims 1-53 stand finally rejected. Appellant filed a Notice of Appeal on August 29, 2005 (as indicated by return of a confirmation postcard marked "OIPE AUG 29 2005"). This brief is submitted in triplicate in support of Appellant's appeal.

TABLE OF CONTENTS

1. REAL PARTY IN INTEREST

The real party in interest is assignee Sendmail, Inc., located at 6425 Christie Ave., 4th Floor, Emeryville, CA 94608.

2. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences known to Appellant, the Appellant's legal representative, or assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

3. STATUS OF CLAIMS

Claims 1-53 are pending in the subject application and are the subject of this appeal. An appendix setting forth the claims involved in the appeal is included as the last section of this brief.

4. STATUS OF AMENDMENTS

Two Amendments have been filed in this case. Appellant mailed an Amendment on April 1, 2005, in response to a non-final Office Action dated December 1, 2004. Net of the Amendment, Appellant believes that the pending claims clearly distinguished the claimed invention over the art of record. In response to the Examiner's Final Rejection dated May 25, 2005, Appellant filed a Notice of Appeal. Appellant has filed an Amendment After Appeal to remove non-art issues (as discussed below in section 6). Based on telephone discussion with the Examiner, it is understood that this Amendment (which makes the single text substitution of "desired" for "user-configurable" in the claims) will be entered. Appellant has chosen to forgo any further amendments that might limit is the scope of Appellant's claims, as it is believed that further amendments to the claims are not warranted in view of the art.

5. SUMMARY OF INVENTION

In accordance with the present invention, operation of an e-mail system is

modified to incorporate a flow control filter (service). During processing of incoming e-mail, each child MTA process (that is created to handle a particular new connection) connects to the flow control filter service, so that it can interact with the service during arrival of a message. (See, e.g., Appellant's Specification, p. 17, lines 20-22) This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). (See, e.g., Appellant's Specification, p. 17, lines 22-26) Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. (See, e.g., Appellant's Specification, p. 17, lines 26-29) Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail. (See, e.g., Appellant's Specification, p. 17, line 29 to p. 18, line 2) Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered. (See, e.g., Appellant's Specification, p. 18, lines 2-5)

The overall methodology of operation may be summarized as follows. The following method steps occur in the context of an incoming message that is being processed by the e-mail system (i.e., MTA forking has already occurred) and now the system is ready to evoke the services of the flow control filter of the present invention. (See, e.g., Appellant's Specification, p. 24, line 36 to p. 25, line 1) Invocation of the flow control filter begins with the MTA (i.e., a child MTA of the original (parent) listener) connecting to the flow control filter (e.g., using Sendmail Milter protocol); the filter accepts the connection. (See, e.g., Appellant's Specification, p. 25, lines 3-5; Fig. 6A at 601) The MTA and the filter perform a handshake sequence, including feature and parameter negotiation. At the conclusion of the handshake sequence, a new thread is

created (i.e., in the flow control engine) for processing the new connection/message. (See, e.g., Appellant's Specification, p. 25, lines 5-9; Fig. 6A at 602) Now, the MTA passes to the filter the corresponding connection information (e.g., IP address and host name) of the sending MTA. (See, e.g., Appellant's Specification, p. 25, lines 10-13; Fig. 6A at 603) Based on the connection information, the filter may look up matching class data from the configuration file. (See, e.g., Appellant's Specification, p. 25, lines 13-14; Fig. 6A at 604) In the event that no matching class data is found, the filter will assume unrestricted access for the host and therefore will accept the connection and message. In that case, the flow control engine thread handling the connection may terminate, as there is no further filtering work to be done for this incoming connection and message; the MTA proceeds normally with no further interaction with the filter. (See, e.g., Appellant's Specification, p. 25, lines 14-19; Fig. 6A at 605) Otherwise, the method proceeds to the following filtering steps. The method tests whether class limits have been reached. (See, e.g., Appellant's Specification, p. 25, line 21; Fig. 6A at 606) In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current connection count. (See, e.g., Appellant's Specification, p. 25, lines 21-23; Fig. 6B at 607) Otherwise (i.e., false), the method terminates with the filter rejecting the connection and returning an administrator-defined error code. (See, e.g., Appellant's Specification, p. 25, lines 23-25; Fig. 6B at 608) In the event that the process did not terminate, the MTA reports the sender information to the filter; this occurs in response to the MAIL FROM SMTP phase. (See, e.g., Appellant's Specification, p. 25, lines 25-27; Fig. 6B at 609)

The method notes the sender (i.e., who is the sender) in the class. The administrator-defined class may include, for example, a sender-based parameter indicating that the filter should note the number of unique senders that have arrived in a given timeframe for this particular host (of the class). (See, e.g., Appellant's Specification, p. 25, line 25 to p. 26, line 3; Fig. 6B at 610) In a manner similar to above, the method tests whether class' sender limits have been reached. (See, e.g., Appellant's Specification, p. 26, lines 3-4; Fig. 6B at 611) In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current

unique sender totals. (See, e.g., Appellant's Specification, p. 26, lines 4-6; Fig. 6B at 612) Otherwise, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 6-8; Fig. 6B at 613) In the event that the filtering process has not terminated based on sender information, the method proceeds to test recipient (RCPT TO) information. The configuration file allows the administrator to define a class that limits the number of unique recipients received for that class, over any given time span. As a given message may have multiple recipients, the step repeats for each recipient (information) of the message. (See, e.g., Appellant's Specification, p. 26, lines 9-13; Fig. 6C at 614) As before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 13-15; Fig. 6C at 615) Otherwise, the method updates the totals and proceeds. (See, e.g., Appellant's Specification, p. 26, line 16; Fig. 6C at 616)

The MTA reports the message body, which may be transmitted as one or more blocks. (See, e.g., Appellant's Specification, p. 26, lines 17-18; Fig. 6C at 617) The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. (See, e.g., Appellant's Specification, p. 26, lines 18-20; Fig. 6C at 618) The MTA reports end of message for the current incoming message. (See, e.g., Appellant's Specification, p. 26, lines 20-21; Fig. 6C at 619) The method compares the message size against class limits specified in the configuration file. (See, e.g., Appellant's Specification, p. 26, lines 21-22; Fig. 6D at 620) Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 22-24; Fig. 6D at 621) Otherwise, the incoming message has passed all filters and is accepted. Now, the method may repeat for other incoming messages. (See, e.g., Appellant's Specification, p. 26, lines 24-26; Fig. 6D at 623)

This approach may be easily scaled, for application on a site-wide basis. In that instance, the flow control filter service monitors the children processes for a number of e-mail servers at a given site. In such a configuration, the flow control filter service would

apply policy on a global (site) basis, instead of on a per server basis. (See, e.g., Appellant's Specification, p. 18, lines 6-9)

6. ISSUES

The issues presented on appeal are:

(1) whether claims 1, 16, 17, 20- 22, 25, 27, 28, 31-34, 41, 47, 48, and 50-53 are unpatentable under 35 U.S.C. 102(e);

(2) whether claims 2-4, 5, 9-11, 18, 19, 23, 24, 26, 35, 37, 38, 42, 44 are unpatentable under 35 U.S.C. 103(a);

(3) whether claims 6, 12, 14, 29, 30, 46 are unpatentable under 35 U.S.C. 103(a); and

(4) whether claims 7, 8, 10, 13, 15, 36, 39, 40, 43, 45, 49 are unpatentable under 35 U.S.C. 103(a).

(Duplication across groups (e.g., Claim 10) is necessitated by the Examiner's rejections.)

Regarding the Examiner's non-art rejections in the Examiner's Final Action Paragraph 1 ("Information Disclosure Statement" rejection) and Paragraph 2 (indefiniteness rejection), it is Appellant's understanding that these rejections are overcome by an Information Disclosure Statement filed October 27, 2005, and by Appellant's Amendment After Appeal filed October 31, 2005.

7. GROUPING OF CLAIMS

For purposes of this appeal, Appellant believes that the following groups of claims are separately patentable under Sections 102 and 103. Thus, the claims do not stand or fall together with respect to the rejections under Sections 102 and 103 but are instead grouped as follows:

Group	Claims
1	1, 16, 17, 20-22, 25, 27, 28, 31-34, 41, 47, 48, 50-53
2	2-4, 5, 9-11, 18, 19, 23, 24, 26, 35, 37, 38, 42, 44
3	6, 12, 14, 29, 30, 46
4	7, 8, 10, 13, 15, 36, 39, 40, 43, 45, 49

(The reasoning supporting separate patentability of the above groups is set forth in detail below, in the Argument section.)

8. ARGUMENT

A. Rejection under Section 102

1. General

A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in the single prior art reference. (See, e.g., MPEP Section 2131.) As will be shown below, the reference fails to teach each and every element set forth in claim 1, as well as other claims of Group 1, and therefore fails to establish anticipation of the claimed invention under Section 102.

2. Group 1 claims: Srivastava et al.

Claims 1, 16-17, 21-22, 25, 27-28, 31-34, 41, 47-48 and 50-53 (and apparently claim 20) stand rejected under 35 U.S.C. 102(e) as being anticipated by Srivastava et al. (U.S. Patent No. 6,374,292), hereinafter referred to "Srivastava"). The Examiner's rejection of claim 1 is representative:

> Regarding claim 1, Srivastava discloses a method for processing an incoming e-mail message that is being received from another domain, the method comprising: receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system; in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis; comparing the request from the particular domain against configurable policy rules; and denying the request if any of said policy rules would be violated. [Srivastava discloses using a process to define a particular domain in an email server. An individual

> domain can be configured to allow all mail to be received if the state of the domain is active (establish a new connection) or if the state of the domain is inactive the domain is suspended from routing mail (denying the request). See Srivastava, column 7, lines 36-59. Srivastava further discusses using a multithreaded process, with each thread handling a connection. Examiner considers this to be equivalent to creating a second process for handling a new connection. Srivastava further states that using a single multithreaded process is beneficial by maximizing performance and stability and by minimizing system resource usage. See Srivastava, column 5, lines 9-15.] By this rationale claim 1 is rejected.

As shown below, Appellant's claimed invention is distinguishable on a variety of grounds.

Srivastava refers to a package of software running on a mail server which governs which specific e-mail related services are offered to users in a particular domain or set of domains. In fact, all of the services described in Srivastava are at the upper OSI layers, i.e., "Presentation" and "Application". Chief among these is virtual domain hosting, whereby one server is given the ability to send and receive e-mail for a variety of otherwise unrelated domains, and to thereby provide e-mail services for users within those domains. The intent of Srivastava is to give the ISP (Internet Service Provider) the means to delegate these services to the administrators of those respective domains without also maintaining individual e-mail servers for each domain or granting machine-wide administrative access to lots of disparate organizations, which of course are prospects with very serious scaling implications.

Appellant's flow control invention, since it operates below the "Presentation" layer in the OSI model, has little knowledge of which domains are stored on the server(s) it protects. It is instead operating in the "Session" layer, with a small amount of information made available to it from lower layers. Moreover it has no knowledge of any protocol or service related to e-mail other than SMTP. It may even not be running on a server providing SMTP service at all. Instead, it has knowledge of the origin of the TCP connection being made to an SMTP service, and makes use of this knowledge to classify the connection and thereby moderate the flow of e-mail into or out of a system. The foregoing is provided to afford the reader context for understanding basic distinctions between the two approaches. While Srivastava's invention is more geared toward

assistance and delegation of administration of e-mail services, Appellant's flow control invention serves to prevent domination of network and system resources by a particular e-mail source. This core distinction will now be reviewed in further detail, with emphasis given to Srivastava's actual teachings, Appellant's specification, and Appellant's specific claim limitations.

Appellant's flow control invention is designed to moderate the flow of e-mail inbound. An example of this might be to limit the maximum number of e-mails, connections over time, or simultaneous connections coming from a **given domain** (say, e.g., "prodigy.net") so that the impact of a sudden surge in spam or a deliberate denial-of-service attack can be detected and quashed without impacting the flow of e-mail from other domains. An outbound policy can also be applied, so that for example a user's desktop infected with a virus that then tries to send e-mail to hundreds or thousands of users is blocked when it is detected. Both of these concepts are known commonly as "traffic shaping", although that term often describes a function implemented by routers at the OSI "Transport" layer and below. This is in contrast to Srivastava, which is essentially an e-mail server provisioning system, allowing a machine or set of machines to be "sliced up" in such a way that many "virtual" domains can be served by a small number of hosts (or one), and the services provided to those domains can be managed by the domain's respective owners.

Turning now to the Examiner's specific basis for rejection, the Examiner analogizes Appellant's flow control filter to Srivastava's virtual domains (Srivastava, Col. 7, lines 36-59).

> Referring now to FIG. 4, showing a flowchart that details a process 500 for defining a virtual domain in accordance with an embodiment of the invention. The process 500 begins at 502 by defining a virtual domain node in the DIT. Once the virtual domain node has been defined, corresponding routing table entries are defined at 504 and at 506, various virtual domain are stored at the virtual domain node. It should be noted that the various virtual domain include a list of services permitted the domain. Such services include IMAP, MAPS, POP3, POP3S, SMTP which in some cases requires presentation of credentials. Other of the services include identification of a domain administrator who is authorized to manage the particular virtual domain

> which includes setting particular user-level for particular users in the domain. These services also include designation of a virtual domain postmaster who identifies email message delivery problems, and a state of the domain.

As clearly shown above, Srivastava at this point is describing the creation of a virtual domain. The purpose of a "virtual domain" is discussed by Srivastava: "It is therefore desirable that an email service provider be able to offer email services to multiple organizations each of which has their own virtual domain and to support the ability to define such domains in the directory and host them on a shared mail server." This is not the same as Appellant's limitation of "said second process being connected to a flow control filter providing filtering on a per-domain basis," which is able to block inbound e-mail traffic -- or allow e-mail traffic -- from different domains, depending on whether a particular given domain is complying with the "configurable policy rules."

For example, if a given domain in Srivastava's system is specified to be a virtual domain for which e-mail services are allowed, then Srivastava's system would permit all e-mail traffic from that virtual domain -- regardless of whether some of that traffic is coming from a user machine engaged in spam or a deliberate denial-of-service attack. There certainly is no mention or passing suggestion in Srivastava that his system also includes some sort of adaptive filter that would then somehow further monitor the virtual domains as they make connections to determine whether the connections for e-mail traffic originating from those domains violate policies in a manner that would cause his system to begin rejecting e-mail traffic until such time as the noncompliant domain returns to compliance. Thus, a detailed review of Srivastava's disclosure that the Examiner relies on for the rejection reveals that it is entirely silent regarding any feature that could function in a manner that is analogous to Appellant's claimed flow control mechanism. In order to achieve Appellant's result (e.g., blocking spam and denial of service attacks) in Srivastava's system, one would have to add Appellant's filtering mechanism to Srivastava's system.

Further, the Examiner points to Srivastava's Col. 5, lines 9-15, which states:

> In the described embodiment, access to the message store 304 is multithreaded thereby allowing a single process to manage a large number of connections since each connection is handled by a thread. In this way, multithreaded access maximizes both performance and scalability by minimizing the system resources required for the management of each connection.

Here, the Examiner contends that Srivastava's multithreaded access to a single message store is the same as Appellant's approach of spawning a second process for handling a new incoming connection.

"Threads" and "processes" are not the same. A "process" is an executing program or task. A "thread" is a part of a process that can execute independently of other parts; it exists within a process and uses the process' resources. Unlike processes, multiple threads run within the same address space and share their process' data. The concepts of threads and processes are well known and well documented in the technical literature. (A copy of Kalev, Danny, "Processes and Threads," ITWorld.com, February 9, 2001, was attached for the Examiner's convenience to Appellant's first-filed Amendment.) The article discussed threads and processes in the context of the Linux operating system, but the discussed concepts apply equally well to other operating systems (e.g., UNIX, Windows, Macintosh OS X).

Without discussing the other deficiencies of Srivastava at this point (e.g., a single "message store" in Srivastava versus multiple incoming connections from different domains in Appellant's system), it is clear that the section that the Examiner cites in Srivastava discusses the use of multiple threads, not the spawning of additional processes. If anything, Srivastava at this point teaches away from Appellant's claimed approach of multiple processes (not Srivastava's approach of a single process with multiple threads).

Appellant's flow control invention provides a facility for moderating the flow of SMTP traffic (connections, aggregate volume, and unique senders) into a server or set of servers. This feature is brought out in Appellant's claims. For example, Appellant's claim 1 recites:

> receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;
>
> (Emphasis added.)

This is an incoming connection from another domain (particularly, an MTA at another domain), for the purpose of doing an MTA-to-MTA email exchange. This is not an operation for servicing user requests.

Further, claim 1 requires:

> in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;
>
> comparing the request from the particular domain against configurable policy rules; and
>
> denying the request if any of said policy rules would be violated.
>
> (Emphasis added.)

Here, the claim requires that the above-mentioned incoming connection is passed through a **domain-specific** filter. This approach allows Appellant's flow control invention to detect and prevent massive spam from being received on incoming connections of a particular domain. Srivastava's approach of granting user services cannot be morphed into system that prevents incoming massive spam from an MTA of a particular domain.

Srivastava is essentially a method for providing virtual hosting services for e-mail and web pages, with the ability to create virtual users within that context and optionally delegate authority to those users to manage parts of the virtual space so provisioned. On startup, Appellant's flow control invention reads a configuration file and then reacts to SMTP traffic it observes. It has no "user-serviceable parts"; only the e-mail administrator has access to read or change its configuration. Although the two approaches converge insofar as they are both related *generally* to providing e-mail service at ISPs and can do

some amount of per-user validity checking, the convergence ends there. They otherwise operate on different types of data and operate in different manners (and in fact in different layers of the OSI protocol stack). By virtue of these core architectural differences, Appellant's claims include very specific claim limitations (explicitly highlighted above) that are not taught or suggested by Srivastava. It is respectfully submitted that these features, as set forth in Appellant's claims, clearly distinguish over Srivastava.

The other independent claims rejected under Section 102 (i.e., claims 21 and 41) include the above-mentioned distinguishing per-domain filtering or policy enforcement claim limitations, and are therefore believed to be allowable for the reasons stated above. (The dependent claims rejected under Section 102 are believed to be allowable by virtue of depending from the foregoing independent claims.) Accordingly, it is respectfully submitted that the claims distinguish over Srivastava, and that the Examiner's rejection under Section 102 should not be sustained.

B. Rejections under Section 103

1. General

Under Section 103(a), a patent may not be obtained if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. To establish a prima facie case of obviousness under this section, the Examiner must establish: (1) that there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings, (2) that there is a reasonable expectation of success, and (3) that the prior art reference (or references when combined) must teach or suggest all the claim limitations. (See e.g., MPEP 2142). The references cited by the Examiner fail to meet these conditions.

2. Group 2 claims: Srivastava combined (in various permutations) with Apache, Mosberger, Ahmed, Shaw, Rakoshitz, and Bates

Claim 35 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and RFC 821. Claim 26 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Apache HTTP Server Configuration Files ("Apache"). Claims 2 and 42 (and apparently claim 3) stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Mosberger et al. (U.S. Patent 6,438,597, "Mosberger"). Claims 18, 19, 23, and 24 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Ahmed et al. (U.S. Patent No. 6,704,772). Claim 9 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw et al. (U.S. Patent No. 6,282,565, "Shaw"). Claims 11 and 44 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of Sash (U.S. Pub. No. 2003/0167250). Claims 4 and 5 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Rakoshitz et al. (U.S. Patent No. 6,816,903, "Rakoshitz"). Claim 10, 37, and 38 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of RFC 821.

For the rejected claims of this group, the Examiner has essentially repeated his Srivastava rejection (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section), but has combined bits and pieces of other art references in an effort to shore up his base Srivastava rejection. The various combinations fail to establish a prima facie rejection under Section 103, as the combined references fail to teach or suggest all the claim limitations of the claims of this group. Importantly, the deficiencies of the base Srivastava rejection are left wholly unaddressed.

For example, Mosberger describes firewall-like features of controlling connections (e.g., based on domain). However, Mosberger does not provide sufficient teaching to morph Srivastava's virtual domain hosting into an e-mail flow control filter that controls connections based on domain-specific behavior of e-mail traffic. As another example, Rakoshitz describes "select counters for monitoring incoming and outgoing traffic from a link" (Rakoshitz, at Col. 21, lines 2-3). Nowhere does Rakoshitz describe maintaining "a counter indicating how many connections have been granted to the particular domain" (emphasis added), as required by Appellant's claim 4, for example. To

the extent that Rakoshitz teaches a counter, the Rakoshitz counter is one that tracks individual links. In particular, no description is given which teaches or suggests that the individual links traceable to or referencing a particular domain be tracked with a counter. Accordingly, at best, Rakoshitz teaches away from Appellant's domain counter claim limitation.

As yet another example, the Examiner adds the Sash permutation/combination for the additional teaching of "a maximum number of different recipients permitted..." Sash describes an information template and describes limiting a maximum number of recipients that an information template can be sent to (i.e., limit the number of times it can be forwarded to other recipients). Appellant's claim limitation states, "said configurable policy rules specify a maximum number of different recipients **permitted by a given domain** over a user-configurable period of time." (See, e.g., Appellant's claim 11.) This would apply, for instance, in this scenario of e-mail traffic coming from a particular domain (e.g., *advertiser.net*) having an inordinate number of different recipients (say, e.g., > 10 million). Placing a restriction on the number of times that a data object can be forwarded (e.g., Sash's restriction on the number of recipients that Sash's information template can be forwarded to) bears little relevance to Appellant's claim limitation.

Without even getting to the issue of whether there is some suggestion or motivation in these references to make the combination suggested by the Examiner, the numerious art rejection permutations that the Examiner has put together for the claims of this group all fails to meet even the most basic threshold of teaching or suggesting all the claim limitations. Importantly, the incremental art references added by the Examiner to create the four Section 103 rejections for the claims of this group each fail to cure the deficiencies of Srivastava regarding the core elements set forth in Appellant's base claims (which the present claims depend from). It is respectfully submitted that the various rejections of claims of this group have each failed to establish a prima facie case of obviousness under Section 103, and therefore it is respectfully requested that the Examiner's rejections of these claims not be sustained.

3. Group 3 claims: Srivastava combined with Spam!

Claims 6, 12, 14, 29, 30 and 46 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! (Cranor and LaMacchia, Communications of the ACM, August 1998). Here, the Examiner relies on Srivastava as above, and adds Spam! for the proposition that it is obvious to add a spam filter. The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava in its failure to teach or suggest Appellant's per-domain flow control filter (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section). Besides not curing the deficiencies of Srivastava, the Spam! reference, as used by the Examiner for rejecting claims of this group, is particularly deficient and thus warrants additional discussion and consideration.

At the outset, it is important to recognize that Appellant does not claim to be the first to invent a spam filter, and Appellant's invention itself is not a spam filter but instead a flow control filter which operates at e-mail server level (message transport agent or "MTA") to monitor the behavior of different domains that are connecting to send incoming e-mail. For example, a "bad" behavior would include a denial-of-service attack, which of course itself is not "spam" (unsolicited e-mail) as that term is generally understood. And, of course, a "spam" filter (such as discussed by the Spam! reference) would be useless against such non-spam malicious behavior. This core distinction will now be examined in further detail, with particular emphasis on Appellant's claim limitations.

With respect to claim 6, for example, Appellant's claim discusses "comparing the sender information **from the particular domain** against said configurable policy rules." This is not simply rejecting an e-mail piece based on it coming from a specific sender that has been blacklisted. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following teaching from Appellant's specification (at page 17, line 20 to page 18, line 5):

> As described above, with each new connection a child MTA process is created. In accordance with the present invention, each child process connects to the flow control filter service, so that it can interact with the

> service during arrival of a message. This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. **By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail encountered for a particular domain over a given period of time.** Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered.
>
> (Emphasis added.)

As described above, the purpose of examining header information is not to accept or reject a given single piece of e-mail based on spam criteria (e.g., blacklisted or not), but is use in conjunction with Appellant's flow control filter to further characterize the given domain that is being monitored. For example, as specified in the passage above, one of the criteria may be "number of different recipients." This requires the system to look at e-mail header information, but note in particular that the header information is not being examined for purposes of identifying an individual piece of e-mail as spam, but instead is being used to further characterize the current behavior of the given **domain** that is being monitored (in order to determine whether server-level intervention is warranted). Should anyone have any doubt as to the intention of Appellant's claim language, one need only refer back to parent claim 1, which clearly describes Appellant's invention as a "flow control filter providing filtering on a **per-domain** basis."

Claim 12 similarly establishes that Appellant's claim limitations are not simply claiming or reading on a spam feature. Appellant's claim 12 discusses comparing the e-

mail message body data against policy rules for a particular domain. This is not simply rejecting an e-mail piece based on it having certain content (e.g., explicit content) that is detected and rejected by a spam filter. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following from Appellant's specification (at page 10, line 27 to page 11, line 2):

> The MTA reports the message body, which may be transmitted as one or more blocks. **The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. The MTA reports end of message for the current incoming message. The method compares the message size against class limits specified in the configuration file.** Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code).
>
> (Emphasis added.)

As described above, the message header may be examined, not for the purpose out of accepting or rejecting a given single piece of e-mail based on spam criteria (e.g., offensive content), but is used in conjunction with Appellant's flow control filter to further characterize the given source -- a particular domain (as required by claim 12, and base claim 1) -- that is being monitored. For example, as specified in the passage above, one of the criteria may be "aggregate total bytes received from a given source over a period of time." (See, e.g., Appellant's claim 46: "determining a maximum aggregate volume of e-mail permitted for the particular domain over a user-configurable period of time.") This requires the system to look at the message body, but note in particular that the message body is not being examined for purposes of identifying the e-mail as having spam content, but instead is being used to further characterize the current behavior of the given domain that is being monitored.

Regarding claim 14, for example, the Examiner contends that this is taught by Spam! at page 79, which indicates that ISPs may limit the number of outbound messages each subscriber can send. However, that is not Appellant's claim limitation. Instead,

claim 14 recites (net of amendments) that "said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a user-configurable period of time." As readily apparent from the claim language, the volume of e-mail being regulated is that coming from a given domain, not that coming from an individual user or subscriber. Thus, for example, applying Appellant's invention, the *uspto.gov* e-mail server could be configured to detect an abnormal amount of e-mail coming from *aol.com* (and take appropriate action, accordingly), for example as a result of a denial-of-service attack originating from that domain. Such a result cannot be achieved by simply using a spam filter at the ISP (*aol.com*) for attempting to detect an abnormally high level of e-mail from any given subscriber. And, in the case of a distributed denial-of-service attack, the attack may be distributed over numerous subscriber accounts (e.g., as a result of Trojan/zombie infection), and it may very well be the case that no one subscriber has an abnormal level of outbound messages.

Regarding claims 29 and 30, for example, the Examiner contends that Spam! at p. 79 teaches that limits can be placed on a domain. However, Spam! itself describes placing limits on individual subscriber accounts. Spam! only describes blocking e-mail from a bogus domain. It contains no description of how a spam filter could continually monitor the e-mail traffic behavior of a given domain, and apply configurable policy rules. Again, the characteristics of e-mail traffic coming from a given domain (e.g., *aol.com*) are not the same as the characteristics of e-mail traffic coming from an individual subscriber (e.g., *john.smith@aol.com*).

For the reasons stated above, the references cited by the Examiner fail to teach or suggest all the claim limitations. Accordingly, it is respectfully submitted that the claims of this group distinguish over the references and are patentable under Section 103. Thus, it is respectfully requested that the Examiner's rejection not be sustained.

4. Group 4 claim: Srivastava/Spam! further combined (individually) with Shaw Bates, and RFC 821

Claim 15 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 14 above, and further in view of Shaw. Claims

8, 43, 45, and 49 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of Bates et al. (U.S. Patent No. 6,779,021, "Bates"). Claim 10 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 9 above, and further in view of RFC 821. Claims 7, 13, 36, 39, and 40 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of RFC 821.

For the rejected claims of this group, the Examiner has essentially repeated his Srivastava/Spam! rejection (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section), but has combined other stray pieces of art in an effort to shore up his rejection. The various combinations fail to establish any competent prima facie rejection under Section 103, as the combinations each fail to teach or suggest all the claim limitations of the claims of this group. Importantly, the deficiencies of the base Srivastava/Spam! Section 103 rejection (discussed above) are left wholly unaddressed. Notable deficiencies in the Examiner's reading of the additional references is evident, as will now be described.

For example, the Examiner relies on Shaw for disclosing the claim limitation of "limiting the size of incoming e-mail messages based on a maximum number of bytes." (See, e.g,. Examiner's Action, paragraph 81.) However, this is not what Appellant's claim states. Instead, claim 15 recites: "said **maximum aggregate volume** is based on total byte count of e-mail received from a given domain over a user-configurable period of time." (Emphasis added.) Claim 14 (claim 15's parent) recites: "wherein said configurable policy rules specify a maximum aggregate volume of e-mail **permitted by a given domain** over a user-configurable period of time." (Emphasis added.) Limiting the maximum average volume from a given domain is not the same as limiting the size of a given incoming e-mail message. Accordingly, the cited art does not serve as an appropriate rejection.

As another example, the Bates passage cited by the Examiner comes from the Bates' Background Section where Bates describes basic spam filtering techniques, such as blocking on sender. Appellant's claim limitation, however, requires: "a maximum

number of different senders permitted by **a given domain** over a user-configurable period of time." (Emphasis added.) A review of Bates indicates no such feature described or suggested. Further, to the extent that Bates repeats spam filtering information about blocking without regard to the behavior of a particular domain, Bates teaches away from Appellant's claimed approach.

The hodgepodge of art combinations that the Examiner has strung together for the claims of this group each fails to meet even the most basic threshold of teaching or suggesting all the claim limitations. And the incremental art references added by the Examiner not only fail to cure the deficiencies of Srivastava/Spam!, but they have also been misinterpreted or mischaracterized by the Examiner to such an extent that their purported incremental teaching is in fact not present in the references themselves. It is respectfully submitted that the rejections of claims of this group each fail to establish a prima facie case of obviousness under Section 103, and therefore it is respectfully requested that the Examiner's rejection of these claims not be sustained.

9. CONCLUSION

The present invention greatly improves the ease and efficiency of running an e-mail server by providing a domain-based e-mail filter. It is respectfully submitted that the present invention, as set forth in the pending claims, sets forth a patentable advance over the art.

In view of the above, it is respectfully submitted that the Examiner's rejections under 35 U.S.C. Section 102 and 103 should not be sustained. If needed, Appellant's undersigned attorney can be reached at 408 884 1507. For the fee due for this Appeal Brief, please refer to the attached Fee Transmittal Sheet. This Brief is submitted in triplicate.

Respectfully submitted,

Date: November 29, 2005

Digitally signed by John A. Smart
Date: 2005.11.29 18:56:56 -08'00'

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX

10.APPENDIX OF CLAIMS ON APPEAL

1. (Original) In an electronic mail (e-mail) system, a method for processing an incoming e-mail message that is being received from another domain, the method comprising:

receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;

in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;

comparing the request from the particular domain against configurable policy rules; and

denying the request if any of said policy rules would be violated.

2. (Previously presented) The method of claim 1, wherein said configurable policy rules specify a maximum number of connections permitted by a given domain over a user-configurable period of time.

3. (Previously presented) The method of claim 2, wherein said user-configurable period of time is configurable.

4. (Original) The method of claim 1, further comprising:

if none of said policy rules would be violated, permitting the requested connection and incrementing a counter indicating how many connections have been granted to the particular domain.

5. (Previously presented) The method of claim 4, further comprising:

after passage of the user-configurable period of time, resetting the counter.

6. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving sender information about the particular e-mail message from the particular domain;

comparing the sender information from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

7. (Original) The method of claim 6, wherein said sender information is transmitted during a "MAIL FROM" phase of SMTP (Simple Mail Transport Protocol) processing.

8. (Previously presented) The method of claim 6, wherein said configurable policy rules specify a maximum number of different senders permitted by a given domain over a user-configurable period of time.

9. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving recipient information about the particular e-mail message from the particular domain;

comparing the recipient information from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

10. (Previously presented) The method of claim 9, wherein said recipient information is transmitted during a "RCPT TO" phase of SMTP (Simple Mail Transport Protocol) processing.

11. (Previously presented) The method of claim 9, wherein said configurable

policy rules specify a maximum number of different recipients permitted by a given domain over a user-configurable period of time.

12. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving e-mail message body data about the particular e-mail message from the particular domain;

comparing the e-mail message body data from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

13. (Previously presented) The method of claim 12, wherein said e-mail message body data is transmitted during a "DATA" phase of SMTP (Simple Mail Transport Protocol) processing.

14. (Previously presented) The method of claim 12, wherein said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a user-configurable period of time.

15. (Previously presented) The method of claim 14, wherein said maximum aggregate volume is based on total byte count of e-mail received from a given domain over a user-configurable period of time.

16. (Original) The method of claim 1, wherein said first process comprises a mail transport agent (MTA) process.

17. (Original) The method claim 16, wherein said second process comprises a child mail transport agent (MTA) process.

18. (Original) The method of claim 1, wherein said second process is created from said first process via a forking operation.

19. (Original) The method of claim 18, wherein said second process is initially created as a copy of said first process.

20. (Original) The method of claim 1, further comprising:

creating a multitude of new processes for handling multiple requests to establish new connections, each new process being connected to said flow control filter providing filtering on a per-domain basis.

21. (Original) An electronic mail (e-mail) system providing filtering of incoming e-mail messages on a per-domain basis, the system comprising:

a parent process for receiving requests from different domains to establish new connections for transmitting e-mail messages;

a plurality of child processes for handling the requests to establish new connections and for handling subsequent requests for transmitting e-mail messages;

a set of rules specifying conditions for accepting requests for new connections and for accepting requests for transmitting e-mail messages; and

a flow control filter, in communication with said child processes and said set of rules, providing filtering based on each domain's conformance to said rules.

22. (Original) The system of claim 21, wherein said parent process and said child processes comprise mail transport agent (MTA) processes.

23. (Original) The system claim 21, wherein each said child process is created from the parent process via a forking operation.

24. (Original) The system of claim 21, wherein each said child process is initially created as a copy of said parent process.

25. (Original) The system of claim 21, wherein said set of rules comprises a configurable set of rules.

26. (Original) The system of claim 21, wherein said set of rules comprises a set of rules stored in a text-based configuration file.

27. (Original) The system of claim 21, wherein said set of rules comprises user-created class definitions specifying different classes of domains.

28. (Original) The system of claim 27, wherein each said class definition includes a domain name corresponding to a particular domain that is to be monitored for filtering.

29. (Previously presented) The system of claim 27, wherein each said class definition includes limits that a particular domain must adhere to over a given user-configurable period of time.

30. (Original) The system of claim 29, wherein said limits include selected ones of:

maximum number of different senders,
maximum number of different recipients,
maximum number of connections,
maximum number of envelopes, and
maximum aggregate volume of mail.

31. (Original) The system of claim 21, wherein a given domain is not filtered if a corresponding rule has not been created for that given domain.

32. (Original) The system of claim 21, wherein said flow control filter denies a given domain's request for a new connection if any of said rules would be violated by

granting the request.

33. (Original) The system of claim 21, wherein said requests for transmitting e-mail messages comprise SMTP (Simple Mail Transport Protocol) commands submitted to the e-mail system from different domains.

34. (Original) The system of claim 33, wherein said flow control filter processes said SMTP commands received from different domains to ascertain whether any of said rules would be violated.

35. (Original) The system of claim 34, wherein said SMTP commands include a "MAIL FROM" command specifying sender information for a given e-mail message.

36. (Original) The system of claim 35, wherein said flow control filter examines said sender information to ascertain whether any of said rules would be violated.

37. (Previously presented) The system of claim 34, wherein said SMTP commands include a "RCPT TO" command specifying recipient information for a given e-mail message.

38. (Original) The system of claim 37, wherein said flow control filter examines said recipient information to ascertain whether any of said rules would be violated.

39. (Original) The system of claim 34, wherein said SMTP commands include a "DATA" command specifying e-mail message body data for a given e-mail message.

40. (Original) The system of claim 39, wherein said flow control filter examines said e-mail message body data to ascertain whether any of said rules would be violated.

41. (Original) In an electronic mail (e-mail) system, a method for processing

incoming e-mail messages that are being received from different domains, the method comprising:

receiving requests from different domains to establish new connections for transmitting e-mail messages to the e-mail system;

for each request received in connection with transmitting a given e-mail message, performing substeps of:

identifying a particular domain that has submitted the request,

based on the determined identity of the domain, determining whether the request to establish a new connection can be granted without violating policy rules, and

based on the determined identity of the domain, determining whether subsequent requests to transmit different portions of a given e-mail message can be granted without violating said policy rules.

42. (Previously presented) The method of claim 41, wherein said step of determining whether the request to establish a new connection can be granted includes:

determining a maximum number of connections permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of connections if the request were granted.

43. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum number of different senders permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different senders if the request were granted.

44. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail

message can be granted includes:

determining a maximum number of different recipients permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different recipients if the request were granted.

45. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum number of different e-mail envelopes permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different e-mail envelopes if the request were granted.

46. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum aggregate volume of e-mail permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum aggregate volume of e-mail if the request were granted.

47. (Original) The method of claim 41, further comprising:

if the request to establish a new connection cannot be granted without violating said policy rules, denying the request.

48. (Original) The method of claim 47, further comprising:

returning an error code indicating why the request is denied.

49. (Original) The method of claim 41, further comprising:

if the request to transmit different portions of a given e-mail message cannot be granted without violating said policy rules, denying the request.

50. (Original) The method of claim 41, wherein portions of a given e-mail message include sender information, recipient information, and message body data.

51. (Original) The method of claim 41, wherein said policy rules are configurable.

52. (Original) The method of claim 41, wherein said policy rules comprise user-edited rules created for different domains.

53. (Previously presented) The method of claim 52, wherein each user-edited rule comprises a host class definition specifying a particular domain and corresponding limits to be applied against that domain over a user-configurable period of time.

PATENT
Docket No. SMI/0005.00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Murray Kucherawy	Examiner: Swearingen, Jeffrey R
Serial No.: 09/945,130	Art Unit: 2145
Filed: August 31, 2001	APPEAL BRIEF
For: E-mail System Providing Filtering Methodology on a Per-Domain Basis	

Mail Stop Appeal
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

BRIEF ON BEHALF OF MURRAY KUCHERAWY.

This is an appeal from the Final Rejection mailed May 25, 2005, in which currently-pending claims 1-53 stand finally rejected. Appellant filed a Notice of Appeal on August 29, 2005 (as indicated by return of a confirmation postcard marked "OIPE AUG 29 2005"). This brief is submitted in triplicate in support of Appellant's appeal.

TABLE OF CONTENTS

1. REAL PARTY IN INTEREST

The real party in interest is assignee Sendmail, Inc., located at 6425 Christie Ave., 4th Floor, Emeryville, CA 94608.

2. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences known to Appellant, the Appellant's legal representative, or assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

3. STATUS OF CLAIMS

Claims 1-53 are pending in the subject application and are the subject of this appeal. An appendix setting forth the claims involved in the appeal is included as the last section of this brief.

4. STATUS OF AMENDMENTS

Two Amendments have been filed in this case. Appellant mailed an Amendment on April 1, 2005, in response to a non-final Office Action dated December 1, 2004. Net of the Amendment, Appellant believes that the pending claims clearly distinguished the claimed invention over the art of record. In response to the Examiner's Final Rejection dated May 25, 2005, Appellant filed a Notice of Appeal. Appellant has filed an Amendment After Appeal to remove non-art issues (as discussed below in section 6). Based on telephone discussion with the Examiner, it is understood that this Amendment (which makes the single text substitution of "desired" for "user-configurable" in the claims) will be entered. Appellant has chosen to forgo any further amendments that might limit is the scope of Appellant's claims, as it is believed that further amendments to the claims are not warranted in view of the art.

5. SUMMARY OF INVENTION

In accordance with the present invention, operation of an e-mail system is

modified to incorporate a flow control filter (service). During processing of incoming e-mail, each child MTA process (that is created to handle a particular new connection) connects to the flow control filter service, so that it can interact with the service during arrival of a message. (See, e.g., Appellant's Specification, p. 17, lines 20-22) This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). (See, e.g., Appellant's Specification, p. 17, lines 22-26) Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. (See, e.g., Appellant's Specification, p. 17, lines 26-29) Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail. (See, e.g., Appellant's Specification, p. 17, line 29 to p. 18, line 2) Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered. (See, e.g., Appellant's Specification, p. 18, lines 2-5)

The overall methodology of operation may be summarized as follows. The following method steps occur in the context of an incoming message that is being processed by the e-mail system (i.e., MTA forking has already occurred) and now the system is ready to evoke the services of the flow control filter of the present invention. (See, e.g., Appellant's Specification, p. 24, line 36 to p. 25, line 1) Invocation of the flow control filter begins with the MTA (i.e., a child MTA of the original (parent) listener) connecting to the flow control filter (e.g., using Sendmail Milter protocol); the filter accepts the connection. (See, e.g., Appellant's Specification, p. 25, lines 3-5; Fig. 6A at 601) The MTA and the filter perform a handshake sequence, including feature and parameter negotiation. At the conclusion of the handshake sequence, a new thread is

created (i.e., in the flow control engine) for processing the new connection/message. (See, e.g., Appellant's Specification, p. 25, lines 5-9; Fig. 6A at 602) Now, the MTA passes to the filter the corresponding connection information (e.g., IP address and host name) of the sending MTA. (See, e.g., Appellant's Specification, p. 25, lines 10-13; Fig. 6A at 603) Based on the connection information, the filter may look up matching class data from the configuration file. (See, e.g., Appellant's Specification, p. 25, lines 13-14; Fig. 6A at 604) In the event that no matching class data is found, the filter will assume unrestricted access for the host and therefore will accept the connection and message. In that case, the flow control engine thread handling the connection may terminate, as there is no further filtering work to be done for this incoming connection and message; the MTA proceeds normally with no further interaction with the filter. (See, e.g., Appellant's Specification, p. 25, lines 14-19; Fig. 6A at 605) Otherwise, the method proceeds to the following filtering steps. The method tests whether class limits have been reached. (See, e.g., Appellant's Specification, p. 25, line 21; Fig. 6A at 606) In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current connection count. (See, e.g., Appellant's Specification, p. 25, lines 21-23; Fig. 6B at 607) Otherwise (i.e., false), the method terminates with the filter rejecting the connection and returning an administrator-defined error code. (See, e.g., Appellant's Specification, p. 25, lines 23-25; Fig. 6B at 608) In the event that the process did not terminate, the MTA reports the sender information to the filter; this occurs in response to the MAIL FROM SMTP phase. (See, e.g., Appellant's Specification, p. 25, lines 25-27; Fig. 6B at 609)

The method notes the sender (i.e., who is the sender) in the class. The administrator-defined class may include, for example, a sender-based parameter indicating that the filter should note the number of unique senders that have arrived in a given timeframe for this particular host (of the class). (See, e.g., Appellant's Specification, p. 25, line 25 to p. 26, line 3; Fig. 6B at 610) In a manner similar to above, the method tests whether class' sender limits have been reached. (See, e.g., Appellant's Specification, p. 26, lines 3-4; Fig. 6B at 611) In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current

unique sender totals. (See, e.g., Appellant's Specification, p. 26, lines 4-6; Fig. 6B at 612) Otherwise, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 6-8; Fig. 6B at 613) In the event that the filtering process has not terminated based on sender information, the method proceeds to test recipient (RCPT TO) information. The configuration file allows the administrator to define a class that limits the number of unique recipients received for that class, over any given time span. As a given message may have multiple recipients, the step repeats for each recipient (information) of the message. (See, e.g., Appellant's Specification, p. 26, lines 9-13; Fig. 6C at 614) As before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 13-15; Fig. 6C at 615) Otherwise, the method updates the totals and proceeds. (See, e.g., Appellant's Specification, p. 26, line 16; Fig. 6C at 616)

The MTA reports the message body, which may be transmitted as one or more blocks. (See, e.g., Appellant's Specification, p. 26, lines 17-18; Fig. 6C at 617) The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. (See, e.g., Appellant's Specification, p. 26, lines 18-20; Fig. 6C at 618) The MTA reports end of message for the current incoming message. (See, e.g., Appellant's Specification, p. 26, lines 20-21; Fig. 6C at 619) The method compares the message size against class limits specified in the configuration file. (See, e.g., Appellant's Specification, p. 26, lines 21-22; Fig. 6D at 620) Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). (See, e.g., Appellant's Specification, p. 26, lines 22-24; Fig. 6D at 621) Otherwise, the incoming message has passed all filters and is accepted. Now, the method may repeat for other incoming messages. (See, e.g., Appellant's Specification, p. 26, lines 24-26; Fig. 6D at 623)

This approach may be easily scaled, for application on a site-wide basis. In that instance, the flow control filter service monitors the children processes for a number of e-mail servers at a given site. In such a configuration, the flow control filter service would

apply policy on a global (site) basis, instead of on a per server basis. (See, e.g., Appellant's Specification, p. 18, lines 6-9)

6. ISSUES

The issues presented on appeal are:

(1) whether claims 1, 16, 17, 20- 22, 25, 27, 28, 31-34, 41, 47, 48, and 50-53 are unpatentable under 35 U.S.C. 102(e);

(2) whether claims 2-4, 5, 9-11, 18, 19, 23, 24, 26, 35, 37, 38, 42, 44 are unpatentable under 35 U.S.C. 103(a);

(3) whether claims 6, 12, 14, 29, 30, 46 are unpatentable under 35 U.S.C. 103(a); and

(4) whether claims 7, 8, 10, 13, 15, 36, 39, 40, 43, 45, 49 are unpatentable under 35 U.S.C. 103(a).

(Duplication across groups (e.g., Claim 10) is necessitated by the Examiner's rejections.)

Regarding the Examiner's non-art rejections in the Examiner's Final Action Paragraph 1 ("Information Disclosure Statement" rejection) and Paragraph 2 (indefiniteness rejection), it is Appellant's understanding that these rejections are overcome by an Information Disclosure Statement filed October 27, 2005, and by Appellant's Amendment After Appeal filed October 31, 2005.

7. GROUPING OF CLAIMS

For purposes of this appeal, Appellant believes that the following groups of claims are separately patentable under Sections 102 and 103. Thus, the claims do not stand or fall together with respect to the rejections under Sections 102 and 103 but are instead grouped as follows:

Group	Claims
1	1, 16, 17, 20-22, 25, 27, 28, 31-34, 41, 47, 48, 50-53
2	2-4, 5, 9-11, 18, 19, 23, 24, 26, 35, 37, 38, 42, 44
3	6, 12, 14, 29, 30, 46
4	7, 8, 10, 13, 15, 36, 39, 40, 43, 45, 49

(The reasoning supporting separate patentability of the above groups is set forth in detail below, in the Argument section.)

8. ARGUMENT

A. Rejection under Section 102

1. General

A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in the single prior art reference. (See, e.g., MPEP Section 2131.) As will be shown below, the reference fails to teach each and every element set forth in claim 1, as well as other claims of Group 1, and therefore fails to establish anticipation of the claimed invention under Section 102.

2. Group 1 claims: Srivastava et al.

Claims 1, 16-17, 21-22, 25, 27-28, 31-34, 41, 47-48 and 50-53 (and apparently claim 20) stand rejected under 35 U.S.C. 102(e) as being anticipated by Srivastava et al. (U.S. Patent No. 6,374,292), hereinafter referred to "Srivastava"). The Examiner's rejection of claim 1 is representative:

> Regarding claim 1, Srivastava discloses a method for processing an incoming e-mail message that is being received from another domain, the method comprising: receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system; in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis; comparing the request from the particular domain against configurable policy rules; and denying the request if any of said policy rules would be violated. [Srivastava discloses using a process to define a particular domain in an email server. An individual

> domain can be configured to allow all mail to be received if the state of the domain is active (establish a new connection) or if the state of the domain is inactive the domain is suspended from routing mail (denying the request). See Srivastava, column 7, lines 36-59. Srivastava further discusses using a multithreaded process, with each thread handling a connection. Examiner considers this to be equivalent to creating a second process for handling a new connection. Srivastava further states that using a single multithreaded process is beneficial by maximizing performance and stability and by minimizing system resource usage. See Srivastava, column 5, lines 9-15.] By this rationale claim 1 is rejected.

As shown below, Appellant's claimed invention is distinguishable on a variety of grounds.

Srivastava refers to a package of software running on a mail server which governs which specific e-mail related services are offered to users in a particular domain or set of domains. In fact, all of the services described in Srivastava are at the upper OSI layers, i.e., "Presentation" and "Application". Chief among these is virtual domain hosting, whereby one server is given the ability to send and receive e-mail for a variety of otherwise unrelated domains, and to thereby provide e-mail services for users within those domains. The intent of Srivastava is to give the ISP (Internet Service Provider) the means to delegate these services to the administrators of those respective domains without also maintaining individual e-mail servers for each domain or granting machine-wide administrative access to lots of disparate organizations, which of course are prospects with very serious scaling implications.

Appellant's flow control invention, since it operates below the "Presentation" layer in the OSI model, has little knowledge of which domains are stored on the server(s) it protects. It is instead operating in the "Session" layer, with a small amount of information made available to it from lower layers. Moreover it has no knowledge of any protocol or service related to e-mail other than SMTP. It may even not be running on a server providing SMTP service at all. Instead, it has knowledge of the origin of the TCP connection being made to an SMTP service, and makes use of this knowledge to classify the connection and thereby moderate the flow of e-mail into or out of a system. The foregoing is provided to afford the reader context for understanding basic distinctions between the two approaches. While Srivastava's invention is more geared toward

assistance and delegation of administration of e-mail services, Appellant's flow control invention serves to prevent domination of network and system resources by a particular e-mail source. This core distinction will now be reviewed in further detail, with emphasis given to Srivastava's actual teachings, Appellant's specification, and Appellant's specific claim limitations.

Appellant's flow control invention is designed to moderate the flow of e-mail inbound. An example of this might be to limit the maximum number of e-mails, connections over time, or simultaneous connections coming from a **given domain** (say, e.g., "prodigy.net") so that the impact of a sudden surge in spam or a deliberate denial-of-service attack can be detected and quashed without impacting the flow of e-mail from other domains. An outbound policy can also be applied, so that for example a user's desktop infected with a virus that then tries to send e-mail to hundreds or thousands of users is blocked when it is detected. Both of these concepts are known commonly as "traffic shaping", although that term often describes a function implemented by routers at the OSI "Transport" layer and below. This is in contrast to Srivastava, which is essentially an e-mail server provisioning system, allowing a machine or set of machines to be "sliced up" in such a way that many "virtual" domains can be served by a small number of hosts (or one), and the services provided to those domains can be managed by the domain's respective owners.

Turning now to the Examiner's specific basis for rejection, the Examiner analogizes Appellant's flow control filter to Srivastava's virtual domains (Srivastava, Col. 7, lines 36-59).

> Referring now to FIG. 4, showing a flowchart that details a process 500 for defining a virtual domain in accordance with an embodiment of the invention. The process 500 begins at 502 by defining a virtual domain node in the DIT. Once the virtual domain node has been defined, corresponding routing table entries are defined at 504 and at 506, various virtual domain are stored at the virtual domain node. It should be noted that the various virtual domain include a list of services permitted the domain. Such services include IMAP, MAPS, POP3, POP3S, SMTP which in some cases requires presentation of credentials. Other of the services include identification of a domain administrator who is authorized to manage the particular virtual domain

> which includes setting particular user-level for particular users in the domain. These services also include designation of a virtual domain postmaster who identifies email message delivery problems, and a state of the domain.

As clearly shown above, Srivastava at this point is describing the creation of a virtual domain. The purpose of a "virtual domain" is discussed by Srivastava: "It is therefore desirable that an email service provider be able to offer email services to multiple organizations each of which has their own virtual domain and to support the ability to define such domains in the directory and host them on a shared mail server." This is not the same as Appellant's limitation of "said second process being connected to a flow control filter providing filtering on a per-domain basis," which is able to block inbound e-mail traffic -- or allow e-mail traffic -- from different domains, depending on whether a particular given domain is complying with the "configurable policy rules."

For example, if a given domain in Srivastava's system is specified to be a virtual domain for which e-mail services are allowed, then Srivastava's system would permit all e-mail traffic from that virtual domain -- regardless of whether some of that traffic is coming from a user machine engaged in spam or a deliberate denial-of-service attack. There certainly is no mention or passing suggestion in Srivastava that his system also includes some sort of adaptive filter that would then somehow further monitor the virtual domains as they make connections to determine whether the connections for e-mail traffic originating from those domains violate policies in a manner that would cause his system to begin rejecting e-mail traffic until such time as the noncompliant domain returns to compliance. Thus, a detailed review of Srivastava's disclosure that the Examiner relies on for the rejection reveals that it is entirely silent regarding any feature that could function in a manner that is analogous to Appellant's claimed flow control mechanism. In order to achieve Appellant's result (e.g., blocking spam and denial of service attacks) in Srivastava's system, one would have to add Appellant's filtering mechanism to Srivastava's system.

Further, the Examiner points to Srivastava's Col. 5, lines 9-15, which states:

> In the described embodiment, access to the message store 304 is multithreaded thereby allowing a single process to manage a large number of connections since each connection is handled by a thread. In this way, multithreaded access maximizes both performance and scalability by minimizing the system resources required for the management of each connection.

Here, the Examiner contends that Srivastava's multithreaded access to a single message store is the same as Appellant's approach of spawning a second process for handling a new incoming connection.

"Threads" and "processes" are not the same. A "process" is an executing program or task. A "thread" is a part of a process that can execute independently of other parts; it exists within a process and uses the process' resources. Unlike processes, multiple threads run within the same address space and share their process' data. The concepts of threads and processes are well known and well documented in the technical literature. (A copy of Kalev, Danny, "Processes and Threads," ITWorld.com, February 9, 2001, was attached for the Examiner's convenience to Appellant's first-filed Amendment.) The article discussed threads and processes in the context of the Linux operating system, but the discussed concepts apply equally well to other operating systems (e.g., UNIX, Windows, Macintosh OS X).

Without discussing the other deficiencies of Srivastava at this point (e.g., a single "message store" in Srivastava versus multiple incoming connections from different domains in Appellant's system), it is clear that the section that the Examiner cites in Srivastava discusses the use of multiple threads, not the spawning of additional processes. If anything, Srivastava at this point teaches away from Appellant's claimed approach of multiple processes (not Srivastava's approach of a single process with multiple threads).

Appellant's flow control invention provides a facility for moderating the flow of SMTP traffic (connections, aggregate volume, and unique senders) into a server or set of servers. This feature is brought out in Appellant's claims. For example, Appellant's claim 1 recites:

> receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;
>
> (Emphasis added.)

This is an incoming connection from another domain (particularly, an MTA at another domain), for the purpose of doing an MTA-to-MTA email exchange. This is not an operation for servicing user requests.

Further, claim 1 requires:

> in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;
>
> comparing the request from the particular domain against configurable policy rules; and
>
> denying the request if any of said policy rules would be violated.
>
> (Emphasis added.)

Here, the claim requires that the above-mentioned incoming connection is passed through a **domain-specific** filter. This approach allows Appellant's flow control invention to detect and prevent massive spam from being received on incoming connections of a particular domain. Srivastava's approach of granting user services cannot be morphed into system that prevents incoming massive spam from an MTA of a particular domain.

Srivastava is essentially a method for providing virtual hosting services for e-mail and web pages, with the ability to create virtual users within that context and optionally delegate authority to those users to manage parts of the virtual space so provisioned. On startup, Appellant's flow control invention reads a configuration file and then reacts to SMTP traffic it observes. It has no "user-serviceable parts"; only the e-mail administrator has access to read or change its configuration. Although the two approaches converge insofar as they are both related *generally* to providing e-mail service at ISPs and can do

some amount of per-user validity checking, the convergence ends there. They otherwise operate on different types of data and operate in different manners (and in fact in different layers of the OSI protocol stack). By virtue of these core architectural differences, Appellant's claims include very specific claim limitations (explicitly highlighted above) that are not taught or suggested by Srivastava. It is respectfully submitted that these features, as set forth in Appellant's claims, clearly distinguish over Srivastava.

The other independent claims rejected under Section 102 (i.e., claims 21 and 41) include the above-mentioned distinguishing per-domain filtering or policy enforcement claim limitations, and are therefore believed to be allowable for the reasons stated above. (The dependent claims rejected under Section 102 are believed to be allowable by virtue of depending from the foregoing independent claims.) Accordingly, it is respectfully submitted that the claims distinguish over Srivastava, and that the Examiner's rejection under Section 102 should not be sustained.

B. Rejections under Section 103

1. General

Under Section 103(a), a patent may not be obtained if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. To establish a prima facie case of obviousness under this section, the Examiner must establish: (1) that there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings, (2) that there is a reasonable expectation of success, and (3) that the prior art reference (or references when combined) must teach or suggest all the claim limitations. (See e.g., MPEP 2142). The references cited by the Examiner fail to meet these conditions.

2. Group 2 claims: Srivastava combined (in various permutations) with Apache, Mosberger, Ahmed, Shaw, Rakoshitz, and Bates

Claim 35 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and RFC 821. Claim 26 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Apache HTTP Server Configuration Files ("Apache"). Claims 2 and 42 (and apparently claim 3) stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Mosberger et al. (U.S. Patent 6,438,597, "Mosberger"). Claims 18, 19, 23, and 24 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Ahmed et al. (U.S. Patent No. 6,704,772). Claim 9 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw et al. (U.S. Patent No. 6,282,565, "Shaw"). Claims 11 and 44 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of Sash (U.S. Pub. No. 2003/0167250). Claims 4 and 5 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Rakoshitz et al. (U.S. Patent No. 6,816,903, "Rakoshitz"). Claim 10, 37, and 38 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of RFC 821.

For the rejected claims of this group, the Examiner has essentially repeated his Srivastava rejection (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section), but has combined bits and pieces of other art references in an effort to shore up his base Srivastava rejection. The various combinations fail to establish a prima facie rejection under Section 103, as the combined references fail to teach or suggest all the claim limitations of the claims of this group. Importantly, the deficiencies of the base Srivastava rejection are left wholly unaddressed.

For example, Mosberger describes firewall-like features of controlling connections (e.g., based on domain). However, Mosberger does not provide sufficient teaching to morph Srivastava's virtual domain hosting into an e-mail flow control filter that controls connections based on domain-specific behavior of e-mail traffic. As another example, Rakoshitz describes "select counters for monitoring incoming and outgoing traffic from a link" (Rakoshitz, at Col. 21, lines 2-3). Nowhere does Rakoshitz describe maintaining "a counter indicating how many connections have been granted to the particular domain" (emphasis added), as required by Appellant's claim 4, for example. To

the extent that Rakoshitz teaches a counter, the Rakoshitz counter is one that tracks individual links. In particular, no description is given which teaches or suggests that the individual links traceable to or referencing a particular domain be tracked with a counter. Accordingly, at best, Rakoshitz teaches away from Appellant's domain counter claim limitation.

As yet another example, the Examiner adds the Sash permutation/combination for the additional teaching of "a maximum number of different recipients permitted..." Sash describes an information template and describes limiting a maximum number of recipients that an information template can be sent to (i.e., limit the number of times it can be forwarded to other recipients). Appellant's claim limitation states, "said configurable policy rules specify a maximum number of different recipients **permitted by a given domain** over a user-configurable period of time." (See, e.g., Appellant's claim 11.) This would apply, for instance, in this scenario of e-mail traffic coming from a particular domain (e.g., *advertiser.net*) having an inordinate number of different recipients (say, e.g., > 10 million). Placing a restriction on the number of times that a data object can be forwarded (e.g., Sash's restriction on the number of recipients that Sash's information template can be forwarded to) bears little relevance to Appellant's claim limitation.

Without even getting to the issue of whether there is some suggestion or motivation in these references to make the combination suggested by the Examiner, the numerious art rejection permutations that the Examiner has put together for the claims of this group all fails to meet even the most basic threshold of teaching or suggesting all the claim limitations. Importantly, the incremental art references added by the Examiner to create the four Section 103 rejections for the claims of this group each fail to cure the deficiencies of Srivastava regarding the core elements set forth in Appellant's base claims (which the present claims depend from). It is respectfully submitted that the various rejections of claims of this group have each failed to establish a prima facie case of obviousness under Section 103, and therefore it is respectfully requested that the Examiner's rejections of these claims not be sustained.

3. Group 3 claims: Srivastava combined with Spam!

Claims 6, 12, 14, 29, 30 and 46 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! (Cranor and LaMacchia, Communications of the ACM, August 1998). Here, the Examiner relies on Srivastava as above, and adds Spam! for the proposition that it is obvious to add a spam filter. The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava in its failure to teach or suggest Appellant's per-domain flow control filter (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section). Besides not curing the deficiencies of Srivastava, the Spam! reference, as used by the Examiner for rejecting claims of this group, is particularly deficient and thus warrants additional discussion and consideration.

At the outset, it is important to recognize that Appellant does not claim to be the first to invent a spam filter, and Appellant's invention itself is not a spam filter but instead a flow control filter which operates at e-mail server level (message transport agent or "MTA") to monitor the behavior of different domains that are connecting to send incoming e-mail. For example, a "bad" behavior would include a denial-of-service attack, which of course itself is not "spam" (unsolicited e-mail) as that term is generally understood. And, of course, a "spam" filter (such as discussed by the Spam! reference) would be useless against such non-spam malicious behavior. This core distinction will now be examined in further detail, with particular emphasis on Appellant's claim limitations.

With respect to claim 6, for example, Appellant's claim discusses "comparing the sender information **from the particular domain** against said configurable policy rules." This is not simply rejecting an e-mail piece based on it coming from a specific sender that has been blacklisted. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following teaching from Appellant's specification (at page 17, line 20 to page 18, line 5):

> As described above, with each new connection a child MTA process is created. In accordance with the present invention, each child process connects to the flow control filter service, so that it can interact with the

> service during arrival of a message. This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. **By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail encountered for a particular domain over a given period of time.** Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered.
>
> (Emphasis added.)

As described above, the purpose of examining header information is not to accept or reject a given single piece of e-mail based on spam criteria (e.g., blacklisted or not), but is use in conjunction with Appellant's flow control filter to further characterize the given domain that is being monitored. For example, as specified in the passage above, one of the criteria may be "number of different recipients." This requires the system to look at e-mail header information, but note in particular that the header information is not being examined for purposes of identifying an individual piece of e-mail as spam, but instead is being used to further characterize the current behavior of the given **domain** that is being monitored (in order to determine whether server-level intervention is warranted). Should anyone have any doubt as to the intention of Appellant's claim language, one need only refer back to parent claim 1, which clearly describes Appellant's invention as a "flow control filter providing filtering on a **per-domain** basis."

Claim 12 similarly establishes that Appellant's claim limitations are not simply claiming or reading on a spam feature. Appellant's claim 12 discusses comparing the e-

mail message body data against policy rules for a particular domain. This is not simply rejecting an e-mail piece based on it having certain content (e.g., explicit content) that is detected and rejected by a spam filter. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following from Appellant's specification (at page 10, line 27 to page 11, line 2):

> The MTA reports the message body, which may be transmitted as one or more blocks. **The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. The MTA reports end of message for the current incoming message. The method compares the message size against class limits specified in the configuration file.** Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code).
>
> (Emphasis added.)

As described above, the message header may be examined, not for the purpose out of accepting or rejecting a given single piece of e-mail based on spam criteria (e.g., offensive content), but is used in conjunction with Appellant's flow control filter to further characterize the given source -- a particular domain (as required by claim 12, and base claim 1) -- that is being monitored. For example, as specified in the passage above, one of the criteria may be "aggregate total bytes received from a given source over a period of time." (See, e.g., Appellant's claim 46: "determining a maximum aggregate volume of e-mail permitted for the particular domain over a user-configurable period of time.") This requires the system to look at the message body, but note in particular that the message body is not being examined for purposes of identifying the e-mail as having spam content, but instead is being used to further characterize the current behavior of the given domain that is being monitored.

Regarding claim 14, for example, the Examiner contends that this is taught by Spam! at page 79, which indicates that ISPs may limit the number of outbound messages each subscriber can send. However, that is not Appellant's claim limitation. Instead,

claim 14 recites (net of amendments) that "said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a user-configurable period of time." As readily apparent from the claim language, the volume of e-mail being regulated is that coming from a given domain, not that coming from an individual user or subscriber. Thus, for example, applying Appellant's invention, the *uspto.gov* e-mail server could be configured to detect an abnormal amount of e-mail coming from *aol.com* (and take appropriate action, accordingly), for example as a result of a denial-of-service attack originating from that domain. Such a result cannot be achieved by simply using a spam filter at the ISP (*aol.com*) for attempting to detect an abnormally high level of e-mail from any given subscriber. And, in the case of a distributed denial-of-service attack, the attack may be distributed over numerous subscriber accounts (e.g., as a result of Trojan/zombie infection), and it may very well be the case that no one subscriber has an abnormal level of outbound messages.

Regarding claims 29 and 30, for example, the Examiner contends that Spam! at p. 79 teaches that limits can be placed on a domain. However, Spam! itself describes placing limits on individual subscriber accounts. Spam! only describes blocking e-mail from a bogus domain. It contains no description of how a spam filter could continually monitor the e-mail traffic behavior of a given domain, and apply configurable policy rules. Again, the characteristics of e-mail traffic coming from a given domain (e.g., *aol.com*) are not the same as the characteristics of e-mail traffic coming from an individual subscriber (e.g., *john.smith@aol.com*).

For the reasons stated above, the references cited by the Examiner fail to teach or suggest all the claim limitations. Accordingly, it is respectfully submitted that the claims of this group distinguish over the references and are patentable under Section 103. Thus, it is respectfully requested that the Examiner's rejection not be sustained.

4. Group 4 claim: Srivastava/Spam! further combined (individually) with Shaw Bates, and RFC 821

Claim 15 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 14 above, and further in view of Shaw. Claims

8, 43, 45, and 49 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of Bates et al. (U.S. Patent No. 6,779,021, "Bates"). Claim 10 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 9 above, and further in view of RFC 821. Claims 7, 13, 36, 39, and 40 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of RFC 821.

For the rejected claims of this group, the Examiner has essentially repeated his Srivastava/Spam! rejection (as discussed above, Appellant's arguments of which are hereby incorporated by reference into this section), but has combined other stray pieces of art in an effort to shore up his rejection. The various combinations fail to establish any competent prima facie rejection under Section 103, as the combinations each fail to teach or suggest all the claim limitations of the claims of this group. Importantly, the deficiencies of the base Srivastava/Spam! Section 103 rejection (discussed above) are left wholly unaddressed. Notable deficiencies in the Examiner's reading of the additional references is evident, as will now be described.

For example, the Examiner relies on Shaw for disclosing the claim limitation of "limiting the size of incoming e-mail messages based on a maximum number of bytes." (See, e.g,. Examiner's Action, paragraph 81.) However, this is not what Appellant's claim states. Instead, claim 15 recites: "said **maximum aggregate volume** is based on total byte count of e-mail received from a given domain over a user-configurable period of time." (Emphasis added.) Claim 14 (claim 15's parent) recites: "wherein said configurable policy rules specify a maximum aggregate volume of e-mail **permitted by a given domain** over a user-configurable period of time." (Emphasis added.) Limiting the maximum average volume from a given domain is not the same as limiting the size of a given incoming e-mail message. Accordingly, the cited art does not serve as an appropriate rejection.

As another example, the Bates passage cited by the Examiner comes from the Bates' Background Section where Bates describes basic spam filtering techniques, such as blocking on sender. Appellant's claim limitation, however, requires: "a maximum

number of different senders permitted by a given domain over a user-configurable period of time." (Emphasis added.) A review of Bates indicates no such feature described or suggested. Further, to the extent that Bates repeats spam filtering information about blocking without regard to the behavior of a particular domain, Bates teaches away from Appellant's claimed approach.

The hodgepodge of art combinations that the Examiner has strung together for the claims of this group each fails to meet even the most basic threshold of teaching or suggesting all the claim limitations. And the incremental art references added by the Examiner not only fail to cure the deficiencies of Srivastava/Spam!, but they have also been misinterpreted or mischaracterized by the Examiner to such an extent that their purported incremental teaching is in fact not present in the references themselves. It is respectfully submitted that the rejections of claims of this group each fail to establish a prima facie case of obviousness under Section 103, and therefore it is respectfully requested that the Examiner's rejection of these claims not be sustained.

9. CONCLUSION

The present invention greatly improves the ease and efficiency of running an e-mail server by providing a domain-based e-mail filter. It is respectfully submitted that the present invention, as set forth in the pending claims, sets forth a patentable advance over the art.

In view of the above, it is respectfully submitted that the Examiner's rejections under 35 U.S.C. Section 102 and 103 should not be sustained. If needed, Appellant's undersigned attorney can be reached at 408 884 1507. For the fee due for this Appeal Brief, please refer to the attached Fee Transmittal Sheet. This Brief is submitted in triplicate.

Respectfully submitted,

Date: November 29, 2005

Digitally signed by John A. Smart
Date: 2005.11.29 18:56:56 -08'00'

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX

10.APPENDIX OF CLAIMS ON APPEAL

1. (Original) In an electronic mail (e-mail) system, a method for processing an incoming e-mail message that is being received from another domain, the method comprising:

receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;

in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;

comparing the request from the particular domain against configurable policy rules; and

denying the request if any of said policy rules would be violated.

2. (Previously presented) The method of claim 1, wherein said configurable policy rules specify a maximum number of connections permitted by a given domain over a user-configurable period of time.

3. (Previously presented) The method of claim 2, wherein said user-configurable period of time is configurable.

4. (Original) The method of claim 1, further comprising:

if none of said policy rules would be violated, permitting the requested connection and incrementing a counter indicating how many connections have been granted to the particular domain.

5. (Previously presented) The method of claim 4, further comprising:

after passage of the user-configurable period of time, resetting the counter.

6. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving sender information about the particular e-mail message from the particular domain;

comparing the sender information from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

7. (Original) The method of claim 6, wherein said sender information is transmitted during a "MAIL FROM" phase of SMTP (Simple Mail Transport Protocol) processing.

8. (Previously presented) The method of claim 6, wherein said configurable policy rules specify a maximum number of different senders permitted by a given domain over a user-configurable period of time.

9. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving recipient information about the particular e-mail message from the particular domain;

comparing the recipient information from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

10. (Previously presented) The method of claim 9, wherein said recipient information is transmitted during a "RCPT TO" phase of SMTP (Simple Mail Transport Protocol) processing.

11. (Previously presented) The method of claim 9, wherein said configurable

policy rules specify a maximum number of different recipients permitted by a given domain over a user-configurable period of time.

12. (Original) The method of claim 1, further comprising:

permitting the requested connection;

receiving e-mail message body data about the particular e-mail message from the particular domain;

comparing the e-mail message body data from the particular domain against said configurable policy rules; and

blocking receipt of the incoming e-mail message if any of said policy rules would be violated.

13. (Previously presented) The method of claim 12, wherein said e-mail message body data is transmitted during a "DATA" phase of SMTP (Simple Mail Transport Protocol) processing.

14. (Previously presented) The method of claim 12, wherein said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a user-configurable period of time.

15. (Previously presented) The method of claim 14, wherein said maximum aggregate volume is based on total byte count of e-mail received from a given domain over a user-configurable period of time.

16. (Original) The method of claim 1, wherein said first process comprises a mail transport agent (MTA) process.

17. (Original) The method claim 16, wherein said second process comprises a child mail transport agent (MTA) process.

18. (Original) The method of claim 1, wherein said second process is created from said first process via a forking operation.

19. (Original) The method of claim 18, wherein said second process is initially created as a copy of said first process.

20. (Original) The method of claim 1, further comprising:

creating a multitude of new processes for handling multiple requests to establish new connections, each new process being connected to said flow control filter providing filtering on a per-domain basis.

21. (Original) An electronic mail (e-mail) system providing filtering of incoming e-mail messages on a per-domain basis, the system comprising:

a parent process for receiving requests from different domains to establish new connections for transmitting e-mail messages;

a plurality of child processes for handling the requests to establish new connections and for handling subsequent requests for transmitting e-mail messages;

a set of rules specifying conditions for accepting requests for new connections and for accepting requests for transmitting e-mail messages; and

a flow control filter, in communication with said child processes and said set of rules, providing filtering based on each domain's conformance to said rules.

22. (Original) The system of claim 21, wherein said parent process and said child processes comprise mail transport agent (MTA) processes.

23. (Original) The system claim 21, wherein each said child process is created from the parent process via a forking operation.

24. (Original) The system of claim 21, wherein each said child process is initially created as a copy of said parent process.

25. (Original) The system of claim 21, wherein said set of rules comprises a configurable set of rules.

26. (Original) The system of claim 21, wherein said set of rules comprises a set of rules stored in a text-based configuration file.

27. (Original) The system of claim 21, wherein said set of rules comprises user-created class definitions specifying different classes of domains.

28. (Original) The system of claim 27, wherein each said class definition includes a domain name corresponding to a particular domain that is to be monitored for filtering.

29. (Previously presented) The system of claim 27, wherein each said class definition includes limits that a particular domain must adhere to over a given user-configurable period of time.

30. (Original) The system of claim 29, wherein said limits include selected ones of:

maximum number of different senders,
maximum number of different recipients,
maximum number of connections,
maximum number of envelopes, and
maximum aggregate volume of mail.

31. (Original) The system of claim 21, wherein a given domain is not filtered if a corresponding rule has not been created for that given domain.

32. (Original) The system of claim 21, wherein said flow control filter denies a given domain's request for a new connection if any of said rules would be violated by

granting the request.

33. (Original) The system of claim 21, wherein said requests for transmitting e-mail messages comprise SMTP (Simple Mail Transport Protocol) commands submitted to the e-mail system from different domains.

34. (Original) The system of claim 33, wherein said flow control filter processes said SMTP commands received from different domains to ascertain whether any of said rules would be violated.

35. (Original) The system of claim 34, wherein said SMTP commands include a "MAIL FROM" command specifying sender information for a given e-mail message.

36. (Original) The system of claim 35, wherein said flow control filter examines said sender information to ascertain whether any of said rules would be violated.

37. (Previously presented) The system of claim 34, wherein said SMTP commands include a "RCPT TO" command specifying recipient information for a given e-mail message.

38. (Original) The system of claim 37, wherein said flow control filter examines said recipient information to ascertain whether any of said rules would be violated.

39. (Original) The system of claim 34, wherein said SMTP commands include a "DATA" command specifying e-mail message body data for a given e-mail message.

40. (Original) The system of claim 39, wherein said flow control filter examines said e-mail message body data to ascertain whether any of said rules would be violated.

41. (Original) In an electronic mail (e-mail) system, a method for processing

incoming e-mail messages that are being received from different domains, the method comprising:

receiving requests from different domains to establish new connections for transmitting e-mail messages to the e-mail system;

for each request received in connection with transmitting a given e-mail message, performing substeps of:

identifying a particular domain that has submitted the request,

based on the determined identity of the domain, determining whether the request to establish a new connection can be granted without violating policy rules, and

based on the determined identity of the domain, determining whether subsequent requests to transmit different portions of a given e-mail message can be granted without violating said policy rules.

42. (Previously presented) The method of claim 41, wherein said step of determining whether the request to establish a new connection can be granted includes:

determining a maximum number of connections permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of connections if the request were granted.

43. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum number of different senders permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different senders if the request were granted.

44. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail

message can be granted includes:

determining a maximum number of different recipients permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different recipients if the request were granted.

45. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum number of different e-mail envelopes permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum number of different e-mail envelopes if the request were granted.

46. (Previously presented) The method of claim 41, wherein said step of determining whether subsequent requests to transmit different portions of a given e-mail message can be granted includes:

determining a maximum aggregate volume of e-mail permitted for the particular domain over a user-configurable period of time; and

determining whether the particular domain would exceed said maximum aggregate volume of e-mail if the request were granted.

47. (Original) The method of claim 41, further comprising:

if the request to establish a new connection cannot be granted without violating said policy rules, denying the request.

48. (Original) The method of claim 47, further comprising:

returning an error code indicating why the request is denied.

49. (Original) The method of claim 41, further comprising:

if the request to transmit different portions of a given e-mail message cannot be granted without violating said policy rules, denying the request.

50. (Original) The method of claim 41, wherein portions of a given e-mail message include sender information, recipient information, and message body data.

51. (Original) The method of claim 41, wherein said policy rules are configurable.

52. (Original) The method of claim 41, wherein said policy rules comprise user-edited rules created for different domains.

53. (Previously presented) The method of claim 52, wherein each user-edited rule comprises a host class definition specifying a particular domain and corresponding limits to be applied against that domain over a user-configurable period of time.